Proceedings

Tenth Annual

Ground Target Modeling and Validation Conference

August 1999

Hosted by
Signature Research, Inc.

Sponsored by
U.S. Army TARDEC

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1a. REPORT SECURITY CLASSIFICATION
Unclassified

1b. RESTRICTIVE MARKINGS
None

2a. SECURITY CLASSIFICATION AUTHORITY

2b. DECLASSIFICATION / DOWNGRADING SCHEDULE

3. DISTRIBUTION / AVAILABILITY OF REPORT
Unlimited

4. PERFORMING ORGANIZATION REPORT NUMBER(S)

5. MONITORING ORGANIZATION REPORT NUMBER(S)

6a. NAME OF PERFORMING ORGANIZATION
Signature Research, Inc.

6b. OFFICE SYMBOL (If applicable)

7a. NAME OF MONITORING ORGANIZATION
U.S. Army Tank-Automotive Command
Research, Engineering & Development Center

6c. ADDRESS (City, State, and ZIP Code)
P.O. Box 346
Calumet, MI 49913

7b. ADDRESS (City, State, and ZIP Code)
Attn.: AMSTA-TR-R
Warren, MI 48397-5000

8a. NAME OF FUNDING / SPONSORING ORGANIZATION

8b. OFFICE SYMBOL (If applicable)

9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER

8c. ADDRESS (City, State, and ZIP Code)

10. SOURCE OF FUNDING NUMBERS

PROGRAM ELEMENT NO.	PROJECT NO.	TASK NO.	WORK UNIT ACCESSION NO.

11. TITLE (Include Security Classification)
Proceedings of the Tenth Annual Ground Target Modeling and Validation Conference (Unclassified)

12. PERSONAL AUTHOR(S)
William R. Reynolds and Tracy T. Maki

13a. TYPE OF REPORT
Conference Proceedings

13b. TIME COVERED
FROM Aug-99 TO _____

14. DATE OF REPORT (Year, Month, Day)
March-00

15. PAGE COUNT
307

16. SUPPLEMENTARY NOTATION

17. COSATI CODES

FIELD	GROUP	SUB-GROUP

18. SUBJECT TERMS (Continue on reverse if necessary and identify by block number)
Signature, Countermeasures, Military Targets, Modeling Validation, Testing, Infrared, Radar/Millimeter Wave, Acoustic, Seismic, Photometric/Visible, Ultraviolet, Magnetic, Simulation, Background Characterization, and Operational Effectiveness

19. ABSTRACT (Continue on reverse if necessary and identify by block number)

This group of papers represents the material presented at the 1999 Conference on Ground Target Modeling and Validation.

The conference serves as a workshop bringing members of the defense research community together to discuss current developments in ground target modeling and validation. It is open to DoD employees, defense industry contractors, academia, and foreign nations.

The focus of the conference is on scientific and technological advancements relating to the tools and techniques of modeling, and the simulation of ground target signatures and background scenes for infrared or near-infrared, radar/millimeter wave, acoustic, seismic, photometric, ultraviolet, and magnetic sensing.

20. DISTRIBUTION / AVAILABILITY OF ABSTRACT
☒ UNCLASSIFIED/UNLIMITED ☐ SAME AS RPT. ☐ DTIC USERS

21. ABSTRACT SECURITY CLASSIFICATION
Unclassified

22a. NAME OF RESPONSIBLE INDIVIDUAL
Grant Gerhart

22b. TELEPHONE (Include Area Code)
810-574-8634

22c. OFFICE SYMBOL
AMSTA-TR-R

DD Form 1473, JUN 86
Previous editions are obsolete.
SECURITY CLASSIFICATION OF THIS PAGE

NOTICES

Disclaimer

The findings of this report are not to be construed as an official Department of the Army position, unless so designated by other authorized documents.

The citation of commercial products or trade names in this report does not constitute an official endorsement of approval of such products.

Distribution Restriction

ACKNOWLEDGEMENTS

Conference Executive Committee

Dr. Grant Gerhart
U.S. Army TARDEC
Warren, MI

Dr. Jack Parks
U.S. Army TARDEC
Warren, MI

Mr. William Reynolds
Signature Research, Inc.
Calumet, MI

Mr. Dave Thomas
U.S. Army TARDEC
Warren, MI

Dr. James Thompson
U.S. Army TARDEC
Warren, MI

FOREWORD

The annual conference on Ground Target Modeling and Validation is jointly sponsored by the U.S. Tank-Automotive Research, Development & Engineering Center (TARDEC). It is administered through Signature Research, Inc. (SGR). The conference is held on the campus of Michigan Technological University in Houghton, Michigan. The Tenth Annual Conference on Ground Target Modeling and Validation was held 17-19 August, 1999.

The conference is held at the UNCLASSIFIED level and is open to all interested persons. In general, the attendance is made up of technical-level individuals representing industry, defense contractors, and government employees. The focus of the conference is on modeling and simulation of ground target signatures and backgrounds, and the validation of such models. Multi-mode technologies include the infrared or near-infrared, radar/millimeter wave, acoustic, seismic, photometric/visible, ultraviolet, and magnetic sensing.

The 1999 proceedings will be provided to all attendees as part of the symposium registration. Contact Signature Research, Inc. at (906) 337-3360, or **sigres@up.net** for additional copies of these proceedings.

Table of Contents

Tenth Annual
Ground Target Modeling & Validation Conference Proceedings

Session 1: Operations Analysis
Session Chair: Mr. Phillip Janicki

Session 2: Active Electromagnetic Signatures
Session Chair: Mr. Mark Pasik

Session 3: Scene Simulation: Projection & Rendering
Session Chair: Mr. Tom Joyner

Session 4: Enabling Ground Vehicle Technologies
Session Chair: Mr. Dave Thomas

Army Space to Ground
Reconnaissance, Surveillance, and Target Acquisition (RSTA)
Operational Analysis

Raymond L. Livingston, Jr.
Teledyne Brown Engineering
Huntsville, AL 35807

ABSTRACT

The Department of Defense (DoD) has long been dependent on various military support functions enabled by space assets. However, the next generation satellite capabilities may profoundly influence Army operations. Soon-to-be-launched space-based Reconnaissance, Surveillance, and Target Acquisition (RSTA) assets may cause radical thinking about how to take maximum advantage of these new capabilities. In particular, RSTA sensor functions in the form of hyper-spectral imagery (HSI) and synthetic aperture radar /ground moving target indicator (SAR/GMTI), will be tested and evaluated on Warfighter 1 and Discoverer II, respectively, in the next few years. In preparation for the evaluation of these new space assets, the Army is embarking on a series of initiatives to determine the military utility of such assets and how best the Army can employ the data and information produced. The Army is concentrating on the tasking, processing, exploitation, and dissemination (TPED) issues. In fact, the TPED issues are the most critical to the Army since it will not own or directly control either satellite system. Central to this front end work are data and analyses to support development of concepts of operations (CONOPS). Issues such as quantity and quality of data, timeliness of dissemination, and echelon of support are just a few examples requiring analysis and experimentation. In order to address a wide array of military utility and TPED issues the US Army Space and Missile Defense Command (SMDC) has a number of initiatives spanning the functions of modeling and simulation, studies and analysis, and experimentation. This paper discusses some of those initiatives and the approaches taken in addressing the military utility and TPED issues. It closes with some of the approaches in developing the necessary models and simulations to support space based RSTA studies and analysis and experimentation for years to come.

INTRODUCTION

There are several initiatives within the DoD community which will provide next-generation space-based RSTA capabilities. In each case, the Army is not the principal driver, but, as a user of space-based RSTA data, it does have requirements that should be determined and then publicized. In order to assist in requirements generation and rationale, appropriate M&S tools must be developed and used. The Army must expand the horizon beyond the current state of modeling and simulation tools. While much analysis will be done by other Services, the National agencies, contractors, etc., the Army must conduct its own analysis based on its own needs and to verify its own requirements. This paper addresses some of these issues from a top-level perspective. It also addresses some of the issues associated with the initiatives to operationalize space.

SCOPE

This paper has several constraining conditions. First, it considers only issues associated with space-based or high altitude endurance unmanned aerial vehicle (HAE-UAV) systems. For the remainder of this paper the term space-based system will refer to both system categories. Second, the paper focuses on the advanced concepts and requirements (ACR) modeling and simulation domain, and not on the other two domains of training, exercises and military operations (TEMO), and research, development and acquisition (RDA). (It is clearly recognized that the majority of the M&S papers being presented at this conference lie within the RDA domain and fall within the taxonomy of systems engineering and analysis.) And, within the ACR domain, the focus is on constructive simulation requirements with emphasis on operational analysis. The paper also addresses issues of relating space-based RSTA data to ground commander operations. Metrics addressing the value of information are discussed. It addresses the need for metrics to address

RSTA issues such as ways to measure the value of space-based RSTA information to the ground commander. And last, the paper confines itself to RSTA operations and does not concern itself with other space-based functions.

NEAR TERM SPACE-BASED RSTA INITIATIVES

The following systems discussions describe a few of the new space-based (or HAE UAV) RSTA platforms undergoing tests or anticipating launch in the next few years. During times of conflict each of these systems is considered a theater asset.

Space-based Radar
The DISCOVERER II (D2) program, formerly known as STARLITE, is a space-based synthetic aperture radar/ground moving target indicator (SAR/GMTI) demonstration initiative. Of particular interest to this audience is the projected capabilities of the imaging SAR system. The D2 SAR is intended to have three levels of resolution: 0.3, 1 and 3 meters with the high resolution mode being a spotlight "staring" capability. D2 was originally undertaken by the Defense Advanced Research Projects Agency (DARPA), the United States Air Force (USAF), and the National Reconnaissance Office (NRO). Its principal objectives are to evolve national asset space reconnaissance technology enabling deployment of a satellite constellation that addresses theater ground reconnaissance and surveillance needs in near real time. The space-based radar (SBR) objective system, and in turn the D2 Demonstration, is a satellite system whose primary purpose is to provide Space Based support to a CINC, Joint Force Commander (JFC), or Service Component Commander during crisis or hostilities. It will augment the in-theater assets and contribute to the overall situational awareness view in a common operations picture. See Figure 1.
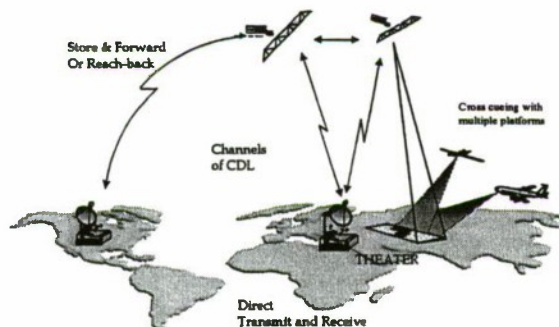
Figure 1. DISCOVERER II

DISCOVERER II is a demonstrator expected to prove a cost-constrained SBR 'objective' system can exist and do the job. The demonstration will baseline direct tasking and product delivery in theater, to maximize responsiveness and to minimize delays. The SBR satellite system will uniquely provide timely, space-based information to users from tactical users to strategic users. The system will use improvements in Tasking, Processing, Exploitation and Dissemination (TPED) processes while integrating with developing Common Data Links (CDL) and Distributed Common Ground Stations (DCGS) for optimizing system's operational use. By 2004, D2's two satellites will demonstrate improved operational capabilities and will provide a path to a financially frugal objective program.

The Army is the current operational customer for D2 through its integration with the TENCAP Tactical Exploitation System (TES) and/or Combined Ground Station (CGS) Module.

The Joint Concept Development Group (JCDG) ensures Navy, Army, and Air Force operational needs are in the D2 CONOPS. The D2 Demonstration will highlight SBR required, and desired, capabilities for future Operational Concept Development.

SMDC is the lead US Army proponent for the DISCOVERER II program. Its mission is to map the Army SAR/GMTI requirements into the design of the DISCOVERER II. The SMDBL is the lead agent for performing the studies, analysis and experimentation incident to resolving the DTLOMS issues.

DISCOVERER II Analysis Issues.
Following are issues related to DISCOVERER II for which operational analyses are appropriate:
- Satellite Direct Tasking and Direct Downlink. The impact of this issue relates to timelines and the availability of data and information. This direct downlink capability could lend itself to receipt and display by tactical maneuver units in near real time.
- Satellite Pass Allocation Strategy. For a joint task force there may be two or more services with multiple echelons vying for D2 data. And, the strategy may change by combat operational phase (i.e., indications & warning, pre-deployment, deployment, operations, redeployment, and peacekeeping). Obviously, there must be some prioritization and deconfliction scheme to address this operational issue. This also has a direct impact on the type and timeliness of data.
- Processing, Exploitation, and Dissemination (PED). This issue addresses timeliness, but also the idea of data adequacy and needs by organizational echelon.
- Quantity, Quality, and Timeliness (QQT). As an overarching term QQT captures the essence of much of what is needed for D2 operational analyses and it is a major operational driver.

Warfighter 1

Warfighter 1 is a Joint commercial and DoD effort involving the use of a hyperspectral imaging (HSI) sensor in space. HSI provides very fine spectral resolution of the composition of material based, in part, by its reflective properties. The USAF is lead for DoD for the 2000 launch with a one-year demonstration and evaluation.

For this HSI satellite evaluation the Army needs a military utility analysis addressing several issues ranging from the effectiveness and utility of the sensor data to operational issues associated with tasking and downlink of data. In the case of the sensor effectiveness the Army must work the data processing, exploitation, and dissemination issues. The operational issues will be similar to those for Discoverer II.

The unique capabilities of an HSI sensor need evaluation. Evaluate across the various combat operational phases: indications & warning, pre-deployment, deployment, operations, redeployment, and peacekeeping.

An early goal of the Army's involvement with Warfighter I is that of developing a space simulation capability with multi-functionality.

Warfighter I Analysis Issues

Following are issues related to Warfighter I for which operational analyses are appropriate:

- Satellite Direct Tasking and Direct Downlink. The impact of this issue relates to timelines and the availability of data and information.
- Satellite Constellation Size & Configuration.
- Satellite Revisit and Coverage Times.
- Processing, Exploitation, and Dissemination (PED). This issue addresses timeliness, but also the idea of data adequacy and needs by echelon.
- Quantity, Quality, and Timeliness (QQT). As an overarching term QQT captures the essence of much of what is needed for Warfighter I operational analyses and it is a major operational driver.

Global Hawk

Global Hawk is the aircraft element of the DoD Tier II Plus reconnaissance system. Global Hawk (See Figure 2)
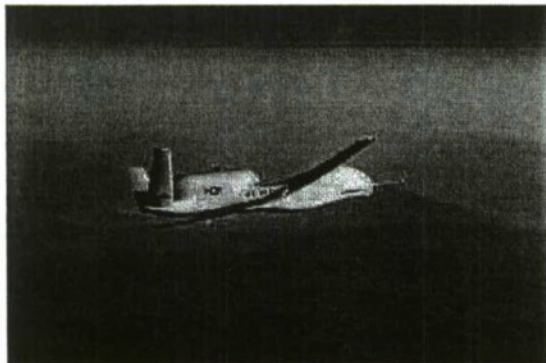
Figure 2. Global Hawk

is a USAF-managed reconnaissance system. It is an attempt to replace expensive manned reconnaissance aircraft such as the U2 and the SR 71, with autonomous, long-endurance unmanned aircraft. Global Hawk has a range capability of 14,000 nautical miles and an endurance of 42 hours. It flies at an altitude of up to 65,000 feet. The aircraft is equipped with optical and infrared cameras and with synthetic aperture radar for continuous surveillance of an area as large as the state of Illinois. The SAR resolution capabilities are expected to be a 1 foot, spotlight mode, and 1 meter in scan mode. It will provide battlespace commanders near-real-time intelligence imagery from high altitudes for long periods of time using SAR, moving target indicator, EO, and infrared sensor systems. A wide-band UHF satellite data link allows direct transmission of imagery data to users in real-time. A Mission Control Element ground station controls mission operations and processes and distributes reconnaissance imagery. Global Hawk prototypes have been built and are currently undergoing flight testing.

Global Hawk Analysis Issues

Following are issues related to Global Hawk for which operational analyses are appropriate:

- System Direct Tasking and Direct Downlink. The impact of this issue relates to timelines and the availability of data and information.
- System Revisit and Coverage Times.
- Processing, Exploitation, and Dissemination (PED). This issue addresses timeliness, but also the idea of data adequacy and needs by echelon.
- Quantity, Quality, and Timeliness (QQT). As an overarching term QQT captures the essence of much of what is needed for Global Hawk operational analyses and it is a major operational driver.

OPERATIONAL ANALYSIS

What is meant by the term operational analysis (OA) as opposed to what might be termed systems analysis, component level analysis, engineering analysis, etc? Unlike some of the work conducted in those latter categories, the OA has a different focus – an operational focus. Other terms for the OA could be a cost and operational effectiveness analysis (COEA), analysis of alternatives (AoA), or military utility analysis. The typical OA has these characteristics:

- Force level focus
- Opposing, force-on-force combat
- Combined arms, and perhaps joint, focus
- System functional interrelationships
- Time period from short battle to multi-day campaign

The operational analysis will more completely represent combined arms forces in an operational environment. The OA uses such documents as ORDs, CONOPS, OMS/MP, etc., to frame its analysis. So, this paper reflects analysis and analytical needs based on those characteristics and conditions.

ARMY ANALYTICAL ISSUES

For these new space-based RSTA assets there are several issues of critical importance to the Army. Some of the more important issues are highlighted below.

- In-theater direct uplink and downlink
- In-theater dynamic tasking and retasking
- Scheduling of theater assets
- Assured access/availability
- Tactical Exploitation of National Capabilities (TENCAP) vs Army Battle Command System (ABCS) trades
- Near continuous surveillance
- Imagery analysis/interpretation
- Automatic/Aided Target Classification and Recognition
- Availability of data.
- Quality, quantity and timeliness of data
- What data is needed? Where? Classification?
- What type data is to be disseminated to which echelon?
- Fusion and correlation of imagery from different sources.

These are operational issues directly impacting the theater commander and his forces.

METRICS

For this discussion two categories of metrics will be established: force effectiveness metrics, and space-based RSTA metrics. The former category is typical of the tactical metrics with which many should be familiar. These include the typical, force-on-force combat endgame metrics such as sensor acquisition ranges, Blue and Red losses, loss exchange ratio (LER), fractional exchange ratio (FER), mission accomplishment, etc.

Now, for space-based RSTA systems, there are several appropriate metrics. These include the following:
- Sensor coverage (instantaneous or over time)
- Revisit frequency
- Redundancy (other space and terrestrial systems)
- Flexibility for mission tasks, types of coverage, etc.
- Quality of data and information
- Quantity of data and information
- Timeliness of data and information

- Tasking, processing, exploitation, and dissemination (TPED) capabilities and timelines

There are also other characteristics of space-based RSTA assets for which other metrics could be described in support of operational analyses. These include the following:

- On-demand reconnaissance
- Near continuous surveillance
- Broad area coverage
- Capability for narrow-area coverage
- Denied area coverage
- Sensor mode characteristics and flexibility
- All weather capability
- Day/night capability

Following are examples of the types of RSTA information these space-based systems provide:

- Situational awareness (a common operational picture) for Blue and Red forces, for pre- and post-hostilities, as well as during combat operations.
- (Precision) Targeting and aided target recognition
- Rapid acquisition and tracking of mobile, time-critical targets.
- Indications and Warning
- Battlefield ordnance awareness (BOA)
- Counter enemy camouflage concealment and deception (CC&D)
- Battle damage assessment (BDA)
- Intelligence preparation of the battlefield (IPB)
- Support to operations other than war

OPERATIONAL ANALYSIS CATEGORIES

As a mechanism for discussing these space to ground RSTA operational analysis issues, three categories can be used. These are space operational analysis, ground operational analysis, and space to ground operational analysis. Analysis of appropriate issues can be conducted within each separate area, or an integrated approach can be used as in the case of the last category. In fact, it is this last category that is the focus of this paper. Part of the community's current shortcomings in this area is that there is the lack of this integrated, end-to-end, space to ground operational analysis environment. The following sections discuss the two separate pieces and then evaluate the need for the integrated environment.

Space Operational Analysis

There are several parameters which are important in determining conditions and capabilities impacting supported unit operations. Some of these are detailed below:

- Constellation types and configurations. Explicit modeling is necessary to determine coverages, revisit rates, effectiveness of various numbers of systems, types of orbits, etc.

- Constellation and orbital mechanics modeling. Modeling here includes such factors as Kepler's equations of motion, altitudes, semi-major axis and eccentricity, numbers of rings and numbers of satellites per ring, types of orbits, inclination angle, and phase angle between rings.

- Sensor parameters. It is important to model the characteristics of various sensor types, combinations of sensors, and space-based sensor effectiveness through realistic environmental conditions over long ranges. Included are the familiar technologies such as EO, IR, RF, etc. Coupled with those are the next generation remote sensing capabilities using multispectral and hyperspectral technologies. Target and background spectral radiance characteristics are especially important. The ability to model realistic conditions through the atmosphere, ground background clutter, target signatures, shadowing, etc., are all required. The impacts of solar insulation and materiel reflectivity, and the impacts of atmospheric absorption and transmissivity are also necessary as modeling and analysis resolution dictate.

Ground Operational Analysis
For this category one has the typical ground analysis issues, but with inputs collected in part from space. These include the following:

- Army Battle Command System (ABCS) modeling. The existing communications architecture and systems must be considered with regard to transmitting data or information from point to point.

- Ground architecture requirements. The entire ground-supporting environment must be understood with respect to numbers of C^4I systems, locations, performance capabilities, linkages, etc. For space systems support these include the typical tasking, processing, exploitation, and dissemination (TPED) systems, as well as the aforementioned ABCS systems. The area involving TPED opens up myriad operational issues involving who, where, how, when, how often, how long, how detailed, etc.

- The area of camouflage, concealment and deception (CC&D) has the typical implications, except now from a remote sensing perspective. There are active and passive systems and techniques from force and systems aspects that have operational impacts.

- The Doctrine, Training, Leader development, Organizations, Materiel, and Soldiers (DTLOMS) reflect issues relevant throughout any analysis and are assumed to be readily understood in the context of OAs.

Space to Ground Operational Analysis
The merging of the space and ground segments produces an integrated environment for conducting space to ground operational analysis. The analysis issues as discussed in the two previous categories also apply here. Additionally, there are other considerations.

- In order to address adequately the combined space and ground elements and functions, an integrated modeling capability is required. This could be a single standalone simulation, or it could be a federation of one or more tools based on analysis requirements.

- An integrated capability provides the analytical environment for conducting military utility analyses and evaluating concepts of operations (CONOPS), tactics, techniques and procedures (TTPs), and in general determining operational requirements.

DATA

The issue of data raises the typical problems practitioners of this operational analysis face. One relishes the results of tests as sources of data, but there are few tests and experiments in this arena to date. And, some of the test data that have been collected are clamped under a classification blanket. Early on, such as when conducting a military utility analysis, the use of parametric bounding of estimated data profiles often proves very useful and even adequate. This idea of using parametric data may not be very appealing to some, but that approach is still required in many operational analyses, so this aspect deserves attention.

CURRENT INITIATIVES

In the area of modeling of space-based RSTA systems and associated operational analyses, there have been several recent initiatives. Some of these are still ongoing, and many point up the need for modeling and simulation (M&S) improvements in the areas covered in this paper. Most of these are Army initiatives, but some reflect work by the broader DoD community.

- Space and Missile Defense Modeling and Simulation Investment Strategy. This is a Space and Missile

Defense Command (SMDC) initiative led by the Space and Missile Defense Battle Lab (SMDBL). It will result in a multi-year strategy for M&S investments for space and missile defense. This strategy is scheduled for completion by 3Q FY 00.

- Space Mission Area Analysis: a top level SMDBL look at future requirements for space-based systems.

- The Battle Command Reengineering Space Initiative is an ongoing, collaborative SMDBL experimentation effort with the Mounted Maneuver Battle Lab evaluating the benefits of putting space products into battalion and brigade TOCs.

- Space Mix Study. This SMDC effort will define the full range of space capabilities required in the outyears. The output will be a future space systems architecture.

- Military utility analysis initiatives. Several such analyses have been completed or are contemplated. These include the Hyperspectral Imagery (HSI) Utility Analysis, and the planned DISCOVERER II User Utility Analysis.

- Spectral Imagery IPT. This is a SMDC effort involving a wide range of spectral imagery initiatives. These include identification and/or development of appropriate modeling and simulation tools, participation in Army and inter-service experiments and exercises, and, in short, assisting in determining the Army's requirements for space-based spectral imaging.

- TENCAP MUSE. This effort involves the use of the Multiple Unified Simulation Environment (MUSE) to model various HAE-UAV and space-based RSTA systems. MUSE is a virtual simulation that supports analyses, experiments and exercises. This initiative's objective is to demonstrate the military utility of Army TENCAP assets in theater and JTF operations.

These are just a few of the Space-based RSTA related initiatives ongoing within the Army and the DoD arena.

MODELING AND SIMULATION ISSUES

From the broad perspective of modeling and simulation there are also issues impacting this subject of space to ground RSTA operational analysis. Some of these are the typical M&S issues as one faces as the state of the art evolves, but each impacts the ability to conduct OAs. For some there are work-arounds, but those are beyond this paper's scope.

- Ability to link space-based RSTA assets and impacts on the ground commander
- "Generic" capability for any space-based RSTA asset
- Some level of resolution to address TPED issues
- What level of functionality to input to the Joint Warfare System (JWARS), the new Joint model for the ACR domain.
- Ground systems signatures consistent with space-based RSTA systems capabilities.
- Large play boxes
- Large numbers of modeled entities
- Processing algorithms
- Relief or work-arounds to classification restrictions for certain systems
- Integration of space and ground tools
- Usable data bases

SUMMARY

This paper has presented a top-level look at issues associated with space to ground RSTA operational analysis. The paper covered emerging space-based RSTA assets, the need for operational analysis, space-based RSTA characteristics and analysis metrics, current initiatives in this area, and related modeling and simulation issues.

ACKNOWLEDGMENTS

The US Army Space and Missile Defense Battle Lab in Huntsville, AL sponsored this paper. The Battle Lab is involved in various aspects of space to ground RSTA operational (and other) analysis initiatives.

REFERENCES

This paper was written with an understanding of ongoing activities within the Space and Missile Defense Command and at the Space and Missile Defense Battle Lab. References for this paper are unpublished and range from assorted briefings, to analysis plans, to current initiatives relating to developing a space-based RSTA operational analysis capability.

Modeling Survivability and Repairability via Markov Chains

Jack Reed, William Jackson, Daniel Hicks
Survivability Optimization Modeling Team
Survivability Technology Center
Research, Development and Engineering Center
US Army Tank-automotive and Armaments Command
AMSTA-TR-R, MS 263, Warren, MI 48397-5000

ABSTRACT

In a sequence of battles, attrition will eventually deplete the Blue force to such an extent that without repairability of damaged platforms (or replenishment) the fighting effectiveness of the Blue force will be compromised. Increased Blue platform survivability will slow, but not halt, the attrition process. Repairability, on the other hand, may maintain the Blue force level at some fraction, depending on both the survivability and repairability rates, of the initial Blue force.

The departure point for this paper, inspired in part by Parks, [1, Parks], is the following simple question: How many vehicles should Blue equip itself with, if Blue forces and Red forces are fighting a sequence of battles for which (1) Blue's survivability rate in each of the battles is .95, (2) during each battle Blue is able to repair and return to the following battle 20% of its damaged vehicles, and (3) Blue needs 40 functional vehicles in order to fight effectively.

This question is answered first with a simple spreadsheet model. The concept of a particular type of Markov process, that of a Markov chain, is then introduced, and it will be shown that the above question is more easily modeled, thought about, and solved, by modeling it as a Markov chain.

The above simple question gives rise to the sequence of questions: What is the relationship between survivability and repairability if x% of the initial Blue force must remain functional throughout an indefinite sequence of battles. A Markov chain model will be used to provide an answer to this question.

Realistically, however, the comparison of repairability and survivability requires that the model used account for the possibility that in each battle some fraction of Blue's force will be damaged beyond repairability. A Markov model to account for this contingency will be developed and

explained.

INTRODUCTION

As mentioned in the abstract, the first goal of the paper is to build a simple spreadsheet model to answer the question: How many vehicles should Blue equip itself with, if Blue forces and Red forces are fighting a sequence of battles for which (1) Blue's survivability rate in each of the battles is .95, (2) during each battle Blue is able to repair and return to the following battle 20% of its damaged vehicles, and (3) Blue needs 40 functional vehicles in order to fight effectively.

Next, the concept of a particular type of Markov process, that of a Markov chain or homogeneous Markov process, is introduced and used to model and answer the above question. The Markov model will then be used to investigate the relationship between the three parameters: (1) Blue force survivability, (2) Blue force repairability, and (3) the fractional value of the Blue force that needs to be always ready for combat.

Lastly, the above model comparing survivability and repairability is made more realistic by allowing combat vehicles to be damaged beyond repairability during one of the sequence of battles. The implications of this more complex model are then presented.

THE SPREADSHEET MODEL

The question to be answered is: How many vehicles should Blue equip itself with, if Blue forces and Red forces are fighting a sequence of battles for which (1) Blue's survivability rate in each of the battles is .95, (2) during each battle Blue is able to repair and return to the following battle 20% of its damaged vehicles, and (3) Blue needs 40 functional vehicles in order to fight effectively.

The initial method for answering the above question is a

spreadsheet model. In the table or spreadsheet given below, the numbers represent a percentage of Blue vehicles, and each row depicts a single battle. The first column numbers the battles, the 0th row gives the initial distribution of vehicles, and the nth row gives the distribution after the nth battle; the second, (U), column gives the percentage of useable Blue vehicles at the beginning of the battle; the third, (R), column gives the percentage of damaged Blue vehicles at the beginning of the battle; the fourth, (U->U), column gives the percentage of Blue vehicles that have survived the battle, 95% of the U column, because of condition (1) in the above question; the fifth, (U->R), column gives the percentage of Blue vehicles damaged during the battle, which is column U minus column U->U; the sixth, (R->U), column gives the percentage of Blue vehicles repaired during the battle, which is 20% of column R, because of condition (2) in the above question; and the last, seventh, (R->R), column gives the percentage of Blue vehicles that remain damaged throughout the battle, which is column R minus column R->U.

Table 1: The Spreadsheet Model

	U	R	U->U	U->R	R->U	R->R
0	100.00	0.00	95.00	5.00	0.00	0.00
1	95.00	5.00	90.25	4.75	1.00	4.00
2	91.25	8.75	86.69	4.56	1.75	7.00
3	88.44	11.56	84.02	4.42	2.31	9.25
4	86.33	13.67	82.01	4.32	2.73	10.94
5	84.75	15.25	80.51	4.24	3.05	12.20
6	83.56	16.44	79.38	4.18	3.29	13.15
7	82.67	17.33	78.54	4.13	3.47	13.86
8	82.00	18.00	77.90	4.10	3.60	14.40
9	81.50	18.50	77.43	4.08	3.70	14.80
10	81.13	18.87	77.07	4.06	3.77	15.10
11	80.84	19.16	76.80	4.04	3.83	15.32
12	80.63	19.37	76.60	4.03	3.87	15.49
13	80.48	19.52	76.45	4.02	3.90	15.62
14	80.36	19.64	76.34	4.02	3.93	15.71
15	80.27	19.73	76.25	4.01	3.95	15.79
16	80.20	19.80	76.19	4.01	3.96	15.84
17	80.15	19.85	76.14	4.01	3.97	15.88
18	80.11	19.89	76.11	4.01	3.98	15.91
19	80.08	19.92	76.08	4.00	3.98	15.93
20	80.06	19.94	76.06	4.00	3.99	15.95
21	80.05	19.95	76.05	4.00	3.99	15.96

22	80.04	19.96	76.03	4.00	3.99	15.97
23	80.03	19.97	76.03	4.00	3.99	15.98
24	80.02	19.98	76.02	4.00	4.00	15.98
25	80.02	19.98	76.01	4.00	4.00	15.99
26	80.01	19.99	76.01	4.00	4.00	15.99
27	80.01	19.99	76.01	4.00	4.00	15.99
28	80.01	19.99	76.01	4.00	4.00	15.99
29	80.00	20.00	76.00	4.00	4.00	16.00
30	80.00	20.00	76.00	4.00	4.00	16.00

Continued computation shows that the values reached in lines 29 and 30 of the spreadsheet model are an equilibrium state. Thus, 80% of the original force can be expected to be useable in any given battle, while 20% of the initial force will be in need of repair.

To answer the question which began this section, if 40 vehicles are needed for combat duty, and only 80% of the total force can be expected to be ready for combat duty, then the initial force should contain 50 vehicles.

SURVIVABILITY AND REPAIRABILITY (SR)

The goal of this section is to formulate the spreadsheet model of the preceding section as a Markov chain. This model, once formulated, will be called the SR model. First, however, Markov chains need to be briefly explained.

A *finite Markov chain*, M, is a system consisting of a finite number of states, $s_1, ..., s_n$, and probabilities, p_{ij}, where p_{ij} is the probability of moving from state s_i to state s_j at any particular time step. Associated with the finite Markov chain, M, is an $n \times n$ matrix $P = [p_{ij}]$, called the *transition matrix* of the chain. Occasionally, an initial distribution of states, $\pi^0 = (\pi^0_1, ..., \pi^0_n)$ is also considered part of the chain.

Now, the spreadsheet model of the preceding section can be reformulated as a Markov chain, giving rise to the SR model. This reformulation is, in reality, quite easy, easier, in fact, that the spreadsheet model. There are two states: U, for useable platforms, is state 1, s_1, and R, for repairable platforms, is state 2, s_2. Now, $p_{11} = .95$, the probability of survival. It follows that $p_{12} = .05$, since from state 1 the platform must either remain in state 1 or move to state 2. Also, $p_{21} = .20$, the repairability rate, and therefore $p_{22} = .80$. Thus, the transition matrix of the model is:

$$P = \begin{bmatrix} .95 & .05 \\ .20 & .80 \end{bmatrix}$$

If the SR model is started with an initial distribution vector, $\pi^0 = (1, 0)$, or $(100, 0)$ if a percent distribution is used as it was in the spreadsheet model, then $\pi^0 P = \pi^1 = (95, 5)$, which gives the values in the U and R columns of the second row (labeled "1") of the spreadsheet model, and in general $\pi^0 P^n = \pi^n$, which will give the values in the U and R columns of the row labeled "n" in the spreadsheet model.

It is now clear the that stable state that was found in the spreadsheet model is just the eigenvector of P. Furthermore, from the theory of positive (all entries are positive) stochastic (all rows sum to 1) matrices, it's known, [2, Lax], that this eigenvector has positive entries, is associated with an eigenvalue of 1, and is unique.

The question that was asked at the beginning of the section on the spreadsheet model can be restated more generally as follows: Given a survivability rate, s, $0 < s < 1$, and a repairability rate, r, $0 < r < 1$, what fraction f of the initial force will remain useable indefinitely? The answer, as indicated above, is the first component of the eigenvector associated with eigenvalue 1 of the matrix $\begin{bmatrix} s & 1-s \\ r & 1-r \end{bmatrix}$.

This gives $f = \dfrac{r}{1 + r - s}$. The original question had $s = .95$ and $r = .20$ so that $f = .20/(1+.20-.95) = .80$ or 80% as was originally found.

To this point, the interest has been in finding f given s and r, and the SR model has provided an easy way to answer that question. Suppose now, however, that f is given. Is there something that can be said about s and r so that the

fraction of useable vehicles will not fall below f throughout a succession of combat engagements? It turns out that there is.

Result: Let $\pi^0 = (1, 0)$, and recall that $\pi^n = (x_n, 1-x_n)$ will be the distribution of the useable and repairable vehicles after the nth battle. If $s \geq f$ and $r \geq \dfrac{f}{1-f}(1-s)$, then $x_n \geq f$ for every n.

This result is relatively straightforward to establish using mathematical induction.

SURVIVABILITY, REPAIRABILITY, AND ANNIHILATION (SRA)

Of course the above model, whether spreadsheet or SR, is severely unrealistic. In an actual engagement, vehicles are sometimes damaged beyond repairability and this possibility needs to be reflected in the model. Either the spreadsheet or the SR model could be extended to reflect the additional possibility of annihilation, but an examination of both models and what can be gleaned from each type should convince the modeler of the efficacy of the SR model. Thus, it is the SR model that is chosen for extension.

The extension of the SR model to the SRA model begins with the addition of a state, state 3 or A, for annihilation. There also needs to be specified some positive probability, a, of a useable vehicle being annihilated during the course of a battle; thus, $p_{13} = a$, and the transition matrix for the SRA model is:

$$Q = \begin{bmatrix} s & 1-(s+a) & a \\ r & 1-r & 0 \\ 0 & 0 & 1 \end{bmatrix} \qquad (1)$$

The immediate thing to notice regarding this model is that state 3 is an *absorbing* state, that is, it can be entered but it can't be left. Hence, regardless of the initial distribution of vehicles, over time the distribution will tend toward $(0, 0, 1)$. The question to be answered then is how fast can this be expected to happen. That is, what is the expected waiting time until absorption, given that the process starts in state 1 or U?

A direct answer to this question can be obtained by evaluating the expression $\sum\limits_{k=1}^{\infty} k\mathrm{P}(\mathrm{U} \text{-} {>} A \text{ in } k \text{ steps})$, which provides the expected number of steps for a platform to be absorbed after starting in state 1 (U).

The contribution of 1-step paths is easy, since there is a unique one-step path from U to A, namely, U->A with probability a.

There is also a unique path from U to A in two steps, namely U->U->A with probability sa.

There are two 3-step paths from U to A, namely U->U->U->A and U->R->U->A with respective probabilities s^2a and $(1-(s+a))ra$.

There are four 4-step paths from U to A, namely U->U->U->U->A, U->R->U->U->A, U->U->R->U->A, and U->R->R->U->A. The first path has probability s^3a, the last $(1-(s+a))(1-r)ra$, and the middle two each have $(1-(s+a))rsa$.

In general, if there are n k-step paths, $k \geq 2$, then there are $2n$ $(k+1)$-step paths, since from each k-step path, two $(k+1)$-step paths can be created, one from inserting a U after the initial U, and one from inserting a R after the initial U. Thus, since there is a unique 2-step path, we have via mathematical induction that there are 2^{k-2} k-step paths for each $k \geq 2$.

Also, from the probability of a k-step path the probabilities of the two $(k+1)$-step paths it generates can be calculated, but keeping track of everything it order to calculate the expected number of steps to absorption using the formula given above is a daunting task. It's here that the advantage of the SRA model over an extended spreadsheet model becomes abundantly clear.

By Markov theory, [3, Isaacson and Madsen], the matrix

$$N = \begin{bmatrix} n_{11} & n_{12} \\ n_{21} & n_{22} \end{bmatrix} = \left(\begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} - \begin{bmatrix} s & 1-(s+a) \\ r & 1-r \end{bmatrix} \right)^{-1}$$

exists and is called the *fundamental* matrix of the model matrix Q, introduced in (1) above. Furthermore, n_{11} is the expected number of times a platform will be in state U given that it started in state U, and n_{12} is the expected number of times such a platform will visit state R. Thus, $n_{11}+n_{12}$ will be the expected number of steps to annihilation for a platform that begins in state U. n_{21} and n_{22} provides analogous information for a platform that begins in state R.

A numerical example should be helpful. Arbitrarily, let p_{11} = .95 (U->U), p_{13} = .01 (U->A), so that p_{12} = .04 (U->R). Also, let the repair rate be 20%. That is, p_{21} = .2 (R->U). Since p_{23} (R->A) is 0, p_{22} = .8 (R->R). Finally, p_{33} = 1 (A->A) so that $p_{31} = p_{32} = 0$. Thus,

$$Q = \begin{bmatrix} .95 & .04 & .01 \\ .2 & .8 & 0 \\ 0 & 0 & 1 \end{bmatrix}, \text{ and}$$

$$N = \left(\begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} - \begin{bmatrix} .95 & .04 \\ .2 & .8 \end{bmatrix} \right)^{-1} = \begin{bmatrix} 100 & 20 \\ 100 & 25 \end{bmatrix}$$

It follows that for a platform that begins in a useable state, the average number of battles until annihilation is 120, and for 20 of those battles, the platform can be expected to be in the repair shop.

The original question asked of the SRA model: How many steps on average until an initially useable platform is annihilated? has been answered quite handily. Can additional information be gleaned from the model? The answer is yes! (Would the question be asked if the answer were No?) Actually, by a slight modification, the model can be tricked into providing a good deal more information.

SURVIVABILITY, REPAIRABILITY, REPLACEMENT (SRP)

Recall that in the SRA model, state 3 (A) was an absorbing state. In the SRP model, state 3 (P) is no longer an absorbing state, but always moves to state 1 (U). That is, in the SRP model, a platform annihilated in one battle is replaced by a new, useable platform in the next. Thus the

matrix of the SRP model is: is: $T = \begin{bmatrix} s & 1-(s+a) & a \\ r & 1-r & 0 \\ 1 & 0 & 0 \end{bmatrix}$.

The questions to be answered are: What will be the expected replacement rate? and: What fraction of the initial force can be expected to be useable, given that replacements for annihilated platforms are effected.

Since the SRP model has a finite number of aperiodic states (a state is aperiodic if the greatest common divisor of the various possible number of steps it takes to return to the state is 1) all of which can communicate, it is known, [3, Isaacson and Madsen], that a limiting distribution of states exists, and that the limiting distribution will be the eigenvector assocciated with the eigenvalue of 1.

An example should be helpful. Take $s = .95$, $r = .2$, and $a = .01$, just as was done in the numerical example in the SRA model. The transition matrix for the SRP model is then

$$T = \begin{bmatrix} .95 & .04 & .01 \\ .2 & .8 & 0 \\ 1 & 0 & 0 \end{bmatrix}.$$ The eigenvector of T associated

with the eigenvalue 1 is $\begin{pmatrix} \dfrac{100}{121} \end{pmatrix}$ or, approximately (.826

.165 .008). It is worth noting that as n tends to infinity, T^n tends to the matrix

$$\begin{bmatrix} \dfrac{100}{121} & \dfrac{20}{121} & \dfrac{1}{121} \\ \dfrac{100}{121} & \dfrac{20}{121} & \dfrac{1}{121} \\ \dfrac{100}{121} & \dfrac{20}{121} & \dfrac{1}{121} \end{bmatrix}.$$

For example, T^{80} is approximately (three decimal places)

$$\begin{bmatrix} .826 & .165 & .008 \\ .826 & .166 & .008 \\ .827 & .165 & .008 \end{bmatrix}.$$

The two questions asked of the SRP model above can now be answered. The expected replacement rate will be slightly under $\dfrac{83}{100}$ s of one percent, and the fraction of the force

expected to be useable will be slightly over $82\dfrac{1}{2}$ percent.

CONCLUSIONS

Models were constructed (SR, SRA, SRP) that permitted relatively easy answers to questions regarding the connections between survivability, repairability, and replacement over a sequence of battles, but it is the technique underlying these models, Markov chains, that is perhaps of the most interest.

For example, one of the exit criteria for a new platform was that ninety percent of the force should survive a forty-eight hour battle. Markov chains would provide a natural way to model this exit criteria.

In electronic warfare, the damage to components during the course of a battle is a topic of ongoing concern. The state of an electronic warfare device also can be naturally modeled using Markov chains.

REFERENCES

[1]   Parks, Jack G, "Survivability, Lethality, and Battlefield Repairability as Markov Processes," *Proceedings of the Military, Government and Aerospace Simulation*, Edited by Michael J. Chinni, The Society for Computer Simulation International, 1999, pp 56-67.

[2]   Lax, Peter D., *Linear Algebra*, John Wiley & Sons, 1997.

[3]   Isaacson, Dean L. and Richard W. Madsen, *Markov Chains Theory and Applications*, Robert E. Kreiger Publishing Company, 1985.

Development of the TOSOM-Force Model

Alan A. Anderson
Teledyne Brown Engineering
1150 Academy Park Loop, Suite 101
Colorado Springs, Colorado, 80910

ABSTRACT

The Threat Oriented Survivability Optimization Model (TOSOM) has been successfully applied numerous times to support system level survivability suite design decisions. The model has proven useful at the system level, successfully supporting a number of recent system survivability suite development efforts. However, development of truly "optimal" survivability strategies requires an assessment of survivability considerations at both the force level and the system level. The current floor version of TOSOM does not provide the information needed for analyzing survivability strategies at the force level, because necessary data, primarily encounter distribution, is not assessed

To address this shortcoming, a variant of TOSOM was developed making the model useful for generating insights into force level survivability strategies. The modified TOSOM has been named TOSOM-Force and is designed to be a standalone variant of the original TOSOM. Specific modifications include:
- the ability to define a common threat environment for multiple system types
- the ability to define encounter distributions at the force level
- the ability to calculate an expected likelihood of achieving specific levels of force survivability

This paper will describe the TOSOM-Force modifications and methodology as well as provide examples as to how the model can be applied to study force level survivability issues.

INTRODUCTION
FORCE LEVEL SURVIVABILITY DISCUSSION

In the interests of finding the "best" survivability approach, the military value of survivability technologies should be assessed at both the individual system level and at the aggregated force level. The system level assessment naturally tends to focus on specific system level attacks. The force level analysis will of necessity examine a longer period of time and will consist of varied attacks distributed unequally over multiple systems.

Several questions immediately jump out when looking at the survivability of the "force". Intuitively, if each component is made as survivable as possible, then the larger unit will also be survivable (although probably not "optimal"). However, the definition of "survival" seems to change as systems are aggregated into systems of systems. At what point is a "force" not survivable? How do you measure survivability at the force level? Traditionally, force survivability goals have been defined as retaining specified percentages of the original force over specified periods of time. For example; "80 percent or more of system A must survive 72 hours of combat". There is some question as to the utility of this approach where forces are non-homogeneous. If all of the air defense elements have been destroyed the force's survivability may have been severely compromised even if 80% or more currently function. The force's ability to resist further losses (particularly against aircraft) may be very low. Even if the force is homogeneous it is possible that a few key losses will severely impact unit effectiveness. For example, suppose that only 10% of the total tanks in a battalion are lost, say a total of 5 systems. But if those systems include the Battalion Commander and each of the 4 line company commanders, although the bulk of the units systems are operational, unit effectiveness (and future unit survivability) has probably been impacted beyond what the numbers would

suggest. Not all individual systems, even of a particular type, are created equal.

The constraints that drive system level decisions also apply at the system of systems (force) level. A very real question to be addressed is "Given a dollar to spend on force survivability, where do I spend it to get the most benefit?"

But which benefit? Systems saved; lives saved; force performance? If we restrict ourselves to the survivability question - providing the most survivability possible for each dollar, we still have a problem in determining and defining the desired or required end state.

How does the definition of survivability change as systems are collected into larger units or "systems of systems"? For a single system we typically treat survivability as a binary state. The system is either "go" or "no go", the system survives or it does not. When designing survivability suites, we attempt to maximize the probability of the "go" outcome given various system design constraints. Addressing the question of force level survivability is much more complicated. Although the force is a collection of various systems, the loss of a portion of those systems does not necessarily indicate a failure to survive on the part of the force. The force can be viewed as a "system of systems" - as long as the larger "force" system is still capable of performing its required function it survives. On the other hand, if the larger system is no longer capable of performing its function, it has failed to "survive", even if significant portions of its subsystems are still operational. There is also the situation where the force can still perform some of its functions at degraded levels. The reality is that survivability in the force level context is not just white or black, go or no go, but a variety of shades of gray.

The TOSOM model currently helps the survivability analyst explore the problem of system level survivability trades. The model treats the system as a target for a predicted array of threat systems and provides insights into which survivability techniques provide the greatest level of protection against an expected value ap-

plication of that specific threat array. As such the model addresses single system types, and utilizes an expected value assessment (weighted average) of the likelihood of encountering various types of threats and the consequence of that encounter. Countermeasures with their associated burdens and anticipated effectiveness are proposed and an estimate of their impact on system survivability is made. By using the estimated suite burdens (cost, weight, etc.) and the expected performance, the analyst can gain insights into the survivability design problem. Typically such an estimate would only be the first step in the analytical process. Potential solutions must then be evaluated in an operational environment (simulated) to determine if they provide the necessary performance. TOSOM can also be visualized as a resource allocation model. It attempts to allocate limited resources (the allowable "burden" levels) while at the same time satisfying the survivability requirements. Sometimes the probability of survival (Ps) requirement is set at a specific level, for example $Ps \geq .9$, or sometimes the goal may be to simply maximize the Ps, while at the same time remaining within the predefined burden constraints.

That same process can be logically extended into the force level survivability problem, but additional assumptions must be made and some definitions agreed upon. The metrics for success change when moving up to a collection of sometimes dissimilar systems. We still want to maximize survivability (or achieve a certain level), but our definition of success is much less clear. Where before we were only concerned with the one system against a single expected value encounter, the force level problem requires the ability to address a multitude of systems of multiple type and varying densities. Where for the single system we were able to use a single expected value attack, the force level problem must consider an uneven distribution of attacks where some systems may be attacked frequently and others not at all. We still want to allocate survivability resources, but the goals are less defined.

For the sake of this discussion ignore system densities for the moment. Assume the force in question is made up of 2 different systems. System A will represent a tank and system B an Infantry Fighting Vehicle (IFV). A and B have different vulnerabilities and different capabilities. Is it more important to improve A's survivability over B's? Do we want them to have the same level of protection? Is that reasonable or desirable?

One way to approach this problem would be to define the desired end state. Say in terms of percentage of the force components surviving a hypothetical conflict. This could be done with high resolution simulations such as JANUS or CAST-FOREM, but set up and run time requirements can be prohibitive. This approach also requires significant computer and analytical resources, and would still require the utilization of a resource allocation model such as TOSOM in order to bring the burden impact question into the study.

In an effort to begin to address some of the problems of force level survivability analysis the TOSOM tool has been expanded into a tool which has been named TOSOM-Force. TOSOM-Force, like TOSOM is NOT a simulation. Both models are designed to provide decision support information to a survivability analyst. While a large amount of data is produced in both models, the goal of TOSOM-Force, like TOSOM, is insight, not numbers. Both tools are designed to work rapidly on a modern personal computer. No other applications (other than Windows) are required to use the programs, although a spread sheet application such as Excel is useful for examining the results.

To best understand TOSOM-Force, the analyst should have familiarity with the standard TOSOM model. Previous papers have provided TOSOM overviews and copies of these papers are available upon request to the author.

TOSOM-Force has two components. The first is an expansion of the original TOSOM threat tree. This expanded tree is again used to collect data on the expected threat array, but it differs from the original TOSOM threat tree by defining the threat encounter and consequence for up to three different systems against the same threat array. The second component of TOSOM-Force is a calculator used to define force size and frequency of encounter and then calculate the probability tables which predict the likelihood of different percentages of force survival.

THREAT TREES

TOSOM-Force is an expansion of the existing TOSOM decision support tool. TOSOM-Force is designed to aid in survivability allocation assessments at a force level. Where TOSOM requires the analyst to develop a threat tree specific to the weapon system under study, TOSOM-Force requires as its first step, the development of a more generic threat tree. The TOSOM-Force threat tree can best be described as the set of all threats to which any member of the force (made up of up to three different types of blue systems) might be exposed. These threats are placed into a TOSOM tree without encounter or consequence probabilities. The generic tree simply defines the structure of the threat. For each of the blue systems a copy of the generic tree is then weighted with the traditional TOSOM encounter and consequence data. Threats which are unique to one system will probably be "zeroed" out in other systems specific trees.

The generic tree approach is justified when it is assumed that each vehicle type that is a member of the "force" has the potential to encounter the same sorts of threats that other friendly system types can encounter. For example, IFV's and Tanks deployed together with likely both encounter the same threat systems. The vulnerability will vary. The actual encounter likelihood may vary. But it is assumed that the opportunity exists to encounter the same threats.

TOSOM-Force expands TOSOM from collecting the data on one system, to collecting the data on as many as three different system types. This is done without significantly increasing the model total run time as once the data is collected only a

single pass through the E2 module is needed to generate the survivability suite performance data for all three systems.

TOSOM-Force essentially can do up to three different TOSOM analyses at the same time. Of added benefit is that the threats are somewhat "calibrated"- the three blue systems are assessed in a "common" threat environment and the survivability suite results are generated with a single pass of the E2 module.

TOSOM-Force modifies the TOSOM threat tree structure, integrating threat encounter data across the variety of vehicle types (requiring survivability enhancement consideration) in the proposed force. The same threat systems (same tree) are used against each of the different blue systems; however, different probability of encounter and consequence values are allowed as a function of the respective vehicle types.

In TOSOM the threat tree allocates the single encounter likelihood across all threats to the system. In TOSOM-Force the actual threat tree structure used is the same for all the friendly systems being modeled. One tree (similar to how we used the same tree for multiple vignettes in the RST-V data set – same tree, different encounter likelihoods). However, each different type blue system could have different values for likelihood and consequence of an encounter with any specific threat. Figure 1 shows the TOSOM-Force threat input screen.
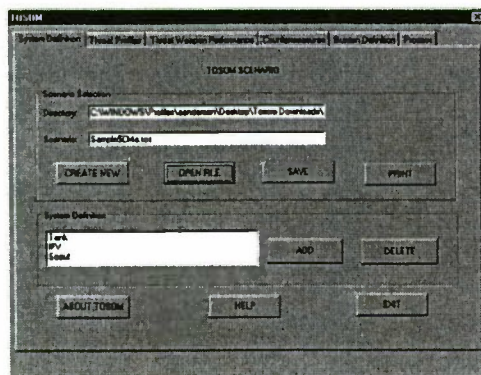
Figure 1.

BURDENS
The burden file has been expanded to specify burden constraints for multiple systems. If the user is studying three different systems (for this sake of this example say a Tank, an IFV and a Scout vehicle), then each of those systems would be expected to (and are allowed to) have different burden budgets.

COUNTERMEASURE DEFINITION
TOSOM-Force allows countermeasures to be designated as system specific or force generic. The E2 module has been modified to allow countermeasures classified as "generic" to be included on any of the force systems when effectiveness calculations are made. System specific countermeasures are allowed only on the designated platforms. An example might help here. A smoke generator would probably be considered to be a "generic" countermeasure. You could probably put the same system on the tank, the IFV and the scout. It would impose the same burdens on each system (cost, weight, etc.). On the other hand, an Armor solution is likely to be unique to the system (system specific) it is designed for. You would not consider putting the tank armor package on the IFV. This approach reduces model run time and data storage requirements and allows the study of larger problems than could be accommodated otherwise. Remember that TOSOM (and TOSOM-Force) use exhaustive enumeration to solve this problem and any approach which reduces the number of unique countermeasures to be considered is useful.

DISCUSSION OF TOSOM E2.
The TOSOM E2 code has been modified to provide survivability and burden calculations for up to three system types included within the force. There are two version of the E2 code. The first version generates all possible combinations of up to 15 defined countermeasures plus the "do nothing option". The E2 code does this by using a binary counting technique. Using binary (base2) arithmetic, it is possible to count to 32,000+ using 16 digits (0's & 1's). Within the E2 code, each countermeasure has an assigned digit (one of the 16 possible). By simply counting from 0 to 32,000+, and examining each binary

number (a 0 in the column means countermeasure not present, and 1 means it is) every possible combination is generated with a single pass (only one loop required). You can think of it as an odometer with 15 "wheels"(or spindles – each representing a different countermeasure) and only 2 digits on each wheel/spindle, a 0 and a 1. Start with all the wheels showing "0"initially (the "do nothing") start "moving"(add "1") and then determine the countermeasure suite that results. By the time all the wheels show "1" (the put everything on it option) you have looked at all possible combinations of 15 countermeasures.

The second version of the E2 generates the combinations for up to 32 different defined countermeasures (under certain restrictions). This is a bit more of a problem. 2^{32} is about 4.2 billion combinations. This would be a significant strain unless we realize that we don't want most of these combinations. By introducing the concept of "feasible" and "acceptable" we can keep both the run time and the required data storage problems under control.

The concept of "feasible" addresses the idea that many of our countermeasures are "either/or" – not both. We may want to consider 5 different armor packages for a vehicle, but only one at a time, not all five (or even 2) at once. This means we don't have to consider any combinations which include more than one of these armor packages. Going back to the odometer example, instead of only a "0" (countermeasure not present) or a "1" (countermeasure present) on the wheel representing our armor countermeasure, lets add a "2", "3", "4", and "5". The "1" represents the presence of Armor option 1, the "2" represents Armor option "2", etc. We can now "spin" the odometer as we did before. Each combination will again be generated, but no combinations which include more than one of the Armor packages will be examined.

The "acceptable" criteria is used to reduce disk storage space. If the burdens resulting for any particular countermeasure suite exceed the user defined burden limits then the data from that calculation are dis-

carded without saving. While this does nothing to relieve the burden on the computer processor, it can significantly relieve the demand for data storage.

FORCE CALCULATOR
Frequency of encounter has still not been addressed. At this point we have TOSOM output similar to what is currently produced, however common threat and countermeasures are built in so that output is automatically calibrated to a baseline threat situation. The same result could be achieved with multiple TOSOM scenarios using the existing model. However data integration problem would be VERY complicated. The way TOSOM-Force addresses the problem of encounter frequency follows.

A data structure for defining encounter frequency distributions was first developed. It might be possible to analyze several standard CASTFOREM scenarios for insights into what these distributions might typically look like. If this analytical approach were to be used the CASTFOREM scenarios could also be used to define the threat trees. This encounter distribution is modeled by apportioning the encounters appropriately among a set of up to 30 vehicles. Figure 2 shows the encounter distribution screen.
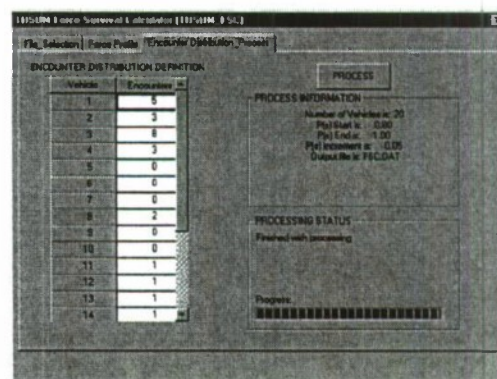

Figure 2

TOSOM-Force uses straight conditional probability to calculate expected levels of force survival. The module output is in a table of the probability of various numbers and combinations of systems

PUTTING IT TOGETHER

The concept is to take the Ps values generated by TOSOM for each potential system survivability suite, and then process each value through an encounter distribution for a fixed quantity of systems (the "Force") involved in a survivability "situation". There are multiple assumptions and simplifications made in order to make this work. The results should still be useful for the purposes of gaining insights into the problem of designing force level survivability countermeasure strategies.

The TOSOM-Force module requires that a probability distribution be provided to model frequency of encounter data. The method used to calculate the probability of different levels of force survival is straightforward, but somewhat tedious – it also has a tendency toward combinatorial explosion for larger sample sizes. The computational intensity can be mitigated by predefinition of the desired force level survivability thresholds, which will be explained later in this paper.

The approach selected calculates the probability of each possible level of surviving force system of a given type. For example, starting with a force of 10 systems, calculate the probability of exactly 0,1,2,...10 systems surviving. The process for doing this is fairly simple, and actually relies upon the same "counting" engine algorithm employed in the exhaustive enumeration (E2) module used in TOSOM. In this case, as in the TOSOM E2 module, there will be 2^n calculations made, where n equals the number of a specific system type within the force (where as in TOSOM n is the number of countermeasures). For a force of 10 tanks, there will be 2^{10} (1,024) calculations required. A force of 20 tanks requires 2^{20} or 1,048,576 calculations.

What TOSOM-Force does is generate every possible combination of surviving members of the "force", one calculation for each possible combination. Each calculation takes into account the analyst provided distribution of encounters. The calculations are then sorted into the appropriate "bins" (0 survivors, 1 survivor,... 10 survivors) and then combined to provide the probability of a specific numbers of

survivors. An example the TOSOM-Force output is provided in table 1 below.

P(s)	x	P(s=x)	P(s>=x)
0.8	0	0	1
0.8	1	0	1
0.8	2	0	1
0.8	3	0	1
0.8	4	0	1
0.8	5	0	1
0.8	6	0	1
0.8	7	0	1
0.8	8	0.000009	1
0.8	9	0.000216	0.999991
0.8	10	0.002249	0.999775
0.8	11	0.013367	0.997526
0.8	12	0.050413	0.984159
0.8	13	0.126541	0.933746
0.8	14	0.21602	0.807205
0.8	15	0.251866	0.591185
0.8	16	0.198306	0.33932
0.8	17	0.102311	0.141014
0.8	18	0.032598	0.038703
0.8	19	0.005699	0.006104
0.8	20	0.000406	0.000406

Table 1.

The chart is interpreted as follows. For the probability distribution used (not shown), the probability of exactly 5 survivors is .142 or about 14%. This number is found by finding the number of survivors in the first column and reading the corresponding value in the 2nd column. The probability of having 5 or more survivors is found by continuing on the corresponding value in the 3rd column which in this case is .332 or about 33%.

CONCLUSION

TOSOM-Force provides a useful tool for investigating some of the issues of interest when studying force level survivability. It helps fill a niche in the fairly sparse array of models and analytical tools available to address the optimal allocation of limited resources for designing or modifying combat systems.

Impact of Aim Point Ambiguity on Vehicle Survivability

Tom Faria and Sharad Kumar

General Dynamics Amphibious Systems

Woodbridge, VA 22191

Abstract

Analysts of direct fire systems have long appreciated the error introduced by a gunner's inability to accurately identify the target center of mass. This component of a detailed error budget is frequently termed the "gun laying error." (Ref. "The Evaluation of Combat Vehicle Fire Control/Gunnery Systems" by K. R. Pfleger and R. J. Bibbero, Report R-1937, dated September 1969) In recent years, this error has been considered as a candidate for improving the survivability of a vehicle: e.g., forcing the threat gunner to increase the "gun laying error" by making it difficult to accurately identify the vehicle "center of mass."

Most combat effectiveness models do not accurately portray this error, but rather attempt to approximate it by modifying the dispersion of the gun system. This note compares the "probability of hit" obtained using the dispersion approach approximation to the "probability of hit" obtained by assuming a distribution of aim points uniformly spread over a region of the target; e.g., an "ambiguous center of mass." The differences in the magnitudes and trends of the two approaches are most dramatic at the shorter ranges, but are still significant at 3000m.

Background

Stealth or signature control technologies have multiple impacts on survivability. The most obvious is, of course, the increase in the time to detect the target vehicle. Of potentially equal importance is the introduction of aim point errors into the fire control error budget of the threat. This is accomplished by misleading the threat gunner as to the target vehicle center of mass; the traditional aim point for most direct fire threats.

Analysts of direct fire systems have long appreciated the impact of not being able to accurately identify the target center of mass (CM). This error is explicitly included in the detailed fire control error budget and is frequently termed the "gun laying error." In the traditional budget, the error results from the target filling a large portion of the sight "field of view" and the gunner not accurately perceiving the location of the vehicle CM. Since the system errors are normally measured in terms of an angle (mr), the "gun laying error" is assumed to decrease with range; e.g., the ambiguity of the fixed size target area presented to the gunner decreases with range in terms of included angle.

It is this aim point ambiguity resulting from the modification of the vehicle signature that is the subject of this note.

The most common approach to accounting for aim point ambiguity is to increase the dispersion of the fire control system while retaining the aim point at the vehicle CM. The dispersion increase is usually accomplished by assuming an additional angular error (mr) which does not decrease with range. This approach of a constant angular error is contrary to the mechanics-physics of the error origin and can lead to an inaccurate assessment the probability of hit and, consequently, survival.

Consider the following. The error can exist only when the threat gunner is aware of the vehicle. Through an undefined cue, he has identified some region on (or adjacent to) the vehicle. This cue is may be a shadow, glint, dust, etc.

19

These cues are equally likely to occur at any position on (or around) the vehicle. The gunner then identifies a region adjacent to or around the cue that does not match the surrounding background. This "region of ambiguity" (with respect to the background) or the cue itself becomes the target. This argument supports the assumption of a uniform distribution of aim points over the presented area of the vehicle rather than an increase in round dispersion about the CM.

The two approaches are depicted in figure 1. The CM of a Bradley has been targeted in the left picture. The inner ellipse represents the dispersions of a baseline threat system. The outer ellipse represents the dispersions increased to account for the lack of a well-defined CM. The right picture is the same Bradley, but with an undefined CM and a random distribution of aim points spread over the presented area. Potentially, even the ground shadow is an aim point.

Common approach   Reality

Increased dispersion to account for aim point ambiguity

Gunner detects target, but cannot identify center of mass

Normal system dispersion
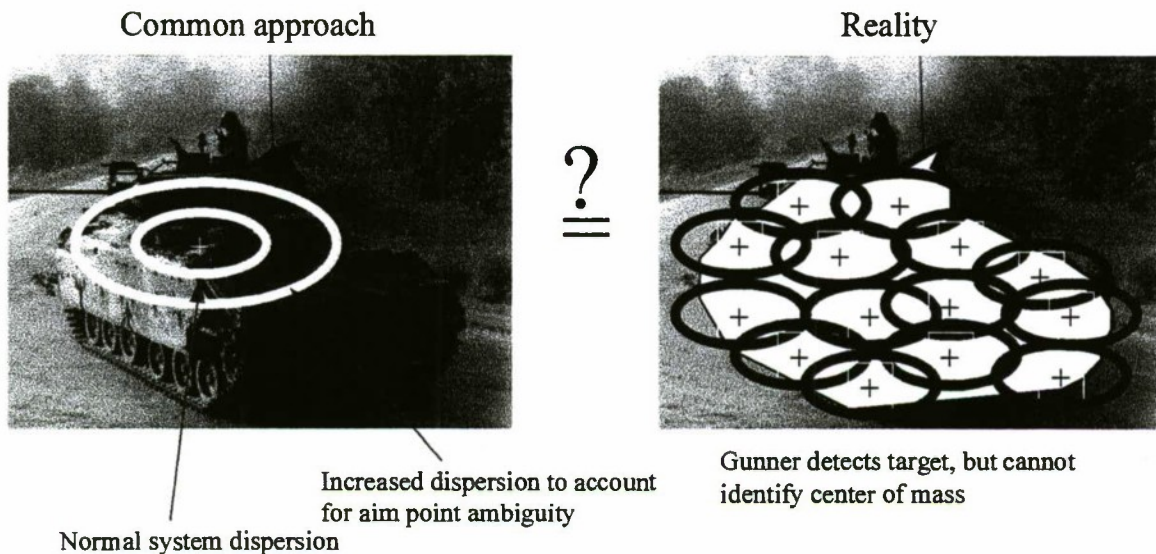
Figure 1: Comparison of Two Approaches

This note quantifies the results of the two approaches, 1) increase in aim point dispersion vice 2) uniform distribution of aim points, for a standard NATO frontal threat gunnery target.

Technical Approach

The technical approach is depicted in figure 2. A standard NATO 2.3m x 2.3m(frontal aspect) target is used for the analysis. The total error budget for the weapon dispersion is assumed to be 0.5 mr for both horizontal (x) and vertical (y) coordinates; these dispersions are typical of a good direct fire medium or large caliber cannon. The probabilities of hit are assumed to follow a Gaussian distribution about the aim point. For the degraded dispersion approach, the target true CM is the aim point for both the baseline (0.5 mr) and degraded (0.6 mr) dispersions. The ambiguous aim point approach assumes that

the aim points are uniformly distributed with respect to the target CM from -X to +X in the horizontal direction and from -Y to +Y in the vertical direction. That is, the "region of ambiguity" for the aim point is +/-X to +/-Y with respect to the target CM. The dispersion for the "region of ambiguity" approach is held constant at 0.5 mr with respect to the actual cannon aim point. The simulation was built in an Excel spreadsheet. The actual aim points for the "region of ambiguity" were drawn randomly from uniform distributions in the horizontal and vertical directions. Trials of 1000 aim point draws (iterations) were averaged for each range and set of conditions. An example of the

20

uniformity of the aim point distribution is shown in figure 3. A plot of the X and Y aim points for one trial (1000 iterations) is shown for a region of ambiguity of -1<=X<=1 and -1<=Y<=1. For a perfectly uniform distribution each of the 20 bins would contain 50 points. The stochastic nature of the simulation is apparent in that number of points per bin ranges from 39 to 65, but uniformity of the X and Y distributions is both apparent and acceptable. The P(hit) results for trials of 1000 iterations are quite repeatable. An example of this repeatability is shown in figure 4 for the "region of ambiguity"

- Standard NATO 2.3m x 2.3m target

- Region of ambiguity defined about target "Center of Mass"

- Actual aimpoints randomly placed (using uniform distribution) within region of ambiguity

- "Probability of Hit" averaged for 1000 iterations
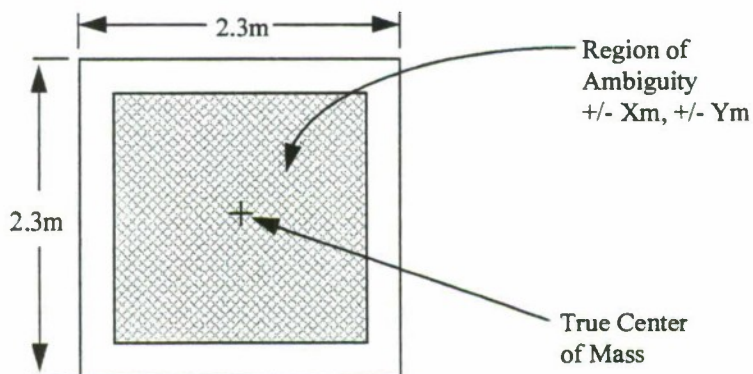
Figure 2: Technical Approach

Aim Point Distributions
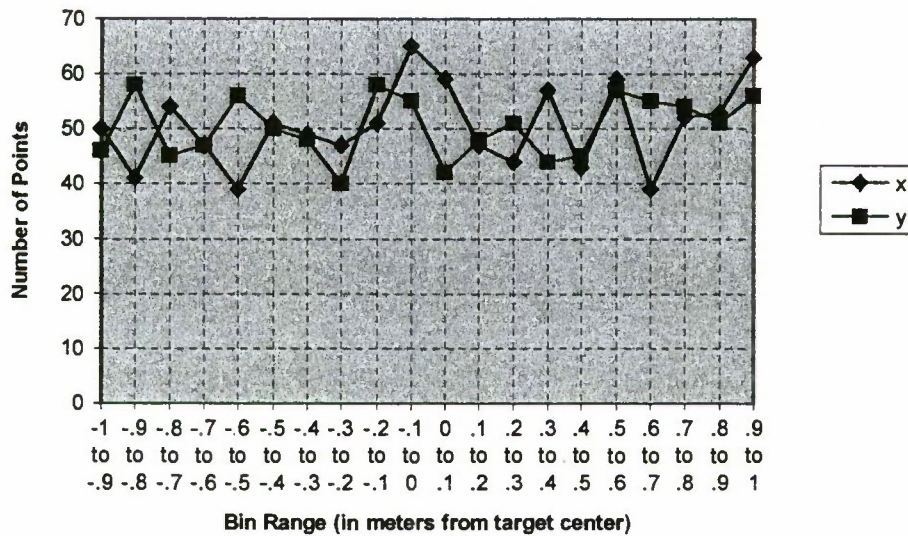
Figure 3: Uniformity of X,Y Aim Point distributions

- Range = 1000m
- X Dispersion = 0.5mr
- Y Dispersion = 0.5mr
- Region of ambiguity
  - -1m < X < 1m
  - -1m < Y < 1m
- 1000 iterations per trial

Trial#	Avg P(hit)
1	0.748
2	0.754
3	0.752
4	0.745
5	0.741
6	0.745
7	0.751
8	0.752
9	0.748
10	0.751
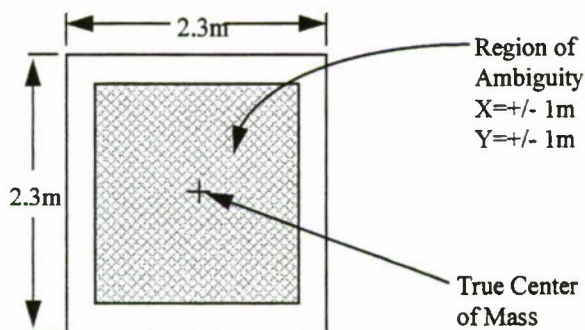
Figure 4: Repeatability of Results

22

case -1<=X<=1 and -1<=Y<=1. P(hit)'s for ten trials of 1000 iterations each are tabulated in the upper right hand corner of the figure. The "probability of hit" for 10 consecutive trials ranged from 0.741 to 0.754. This is a "peak-to-peak" difference of 0.013; e.g., less than 2% of the band. This level of repeatability is adequate for a study of this nature.

Results

The results of this academic exercise are presented in figure 5. "Probabilities of hit" are plotted and tabulated for ranges between 500 meters and 3000 meters at increments of 500 meters. Analysis results are labeled according to the size of the "region of ambiguity" and the weapon dispersion. Dispersions and "regions of ambiguity" dimensions are assumed to be the same in both the horizontal and vertical directions. For example, the label "+/-1m|0.5mr" indicates that the region of ambiguity is -1<=X<=1 and -1<=Y<=1 and the dispersions are 0.5 mr in both the horizontal and vertical directions. A "region of 0 ambiguity" indicates that the gunner can exactly identify the target CM. Results are presented for five sets of conditions "+/-0m|0.5mr," " +/-0m|0.6mr," "+/-0.5m|0.5mr," "+/-1m|0.5mr," and "+/-1.5m|0.5mr." The "region of ambiguity" of one square meter or approximately 60% (1m^2/1.69m^2) of the target area will be used as a discussion example. In the primary ranges of interest, 1km to 2.5km, the results of the two approaches differ in both the magnitude and trend (shape). For this 60% ambiguity size region P(hit)'s for 1000m and 2500m are 0.745 and 0.354 respectively. Increasing the dispersion of the baseline gun by 20% (0.5 mr to 0.6 mr; e.g., adding an additional 0.33mr error) provides the best overall match to a "region of ambiguity" equal to 60% of the total target area. P(hit)'s for a dispersion of 0.6mr at 1000m and 2500m are 0.892 and 0.31 respectively. The greatest difference in P(hit) is at the 1km range; the P(hit) for the increased dispersion is 17% higher than the ambiguous aim point approach. The two approaches cross at approximately 1800m. At 2.5km, the ambiguous aim point approach is 14% higher than the increased dispersion approach (and the "per cent" difference will continue to increase with range). Due to the inherent differences in the two approaches similar results will obtain for all conditions and ranges.
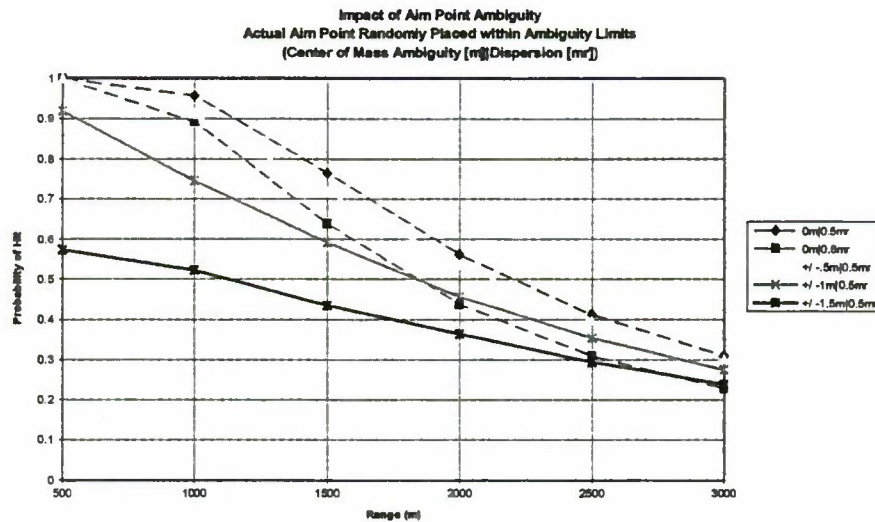
Figure 5: Aim Point Ambiguity Study Results

range(m)	0m|0.5mr	0m|0.6mr	+/ -.5m|0.5mr	+/ -1m|0.5mr	+/ -1.5m|0.5mr
500	1	1	1	0.919	0.574
1000	0.957	0.892	0.911	0.745	0.523
1500	0.765	0.638	0.717	0.592	0.435
2000	0.562	0.438	0.535	0.457	0.364
2500	0.413	0.31	0.397	0.354	0.294
3000	0.31	0.228	0.301	0.275	0.238

Conclusions

The conclusions are summarized in figure 6. In the primary range of interest for current direct fire cannons, 1km to 2.5km, both the magnitudes and the trends of the results for the two approaches differ. The differences are inherent to the mechanics of the approaches and cannot be eliminated by a judicious choice of parameters. The magnitudes of the P(hit) differences can easily exceed 10% This difference will carry through to the probability of survival. Since 10% is frequently considered the decision point for justifying a new system, this error is potentially adequate to cause an unreasonable choice in the final vehicle design.

24

Conclusions

- A 10% change is frequently considered the minimun necessary to justify a new system
- For direct fire cannons, the difference in the "increased dispersion" and "area of ambiguity" approaches can exceed 10%
- Choice of "increased dispersion" approach can lead to a poor choice in system design

Figure 6: Conclusions

Target Acquisition Modeling with Automated Environmental Data Ingest for Weapon System Evaluation

David Sauter and Richard Shirkey
Army Research Laboratory
White Sands Missile Range, NM 88002 USA

ABSTRACT

Target acquisition modeling plays an important role in the mission planning process associated with both man in the loop and autonomous command and control systems. The type of information required typically varies as a function of the amount of time before a mission. In general, early in the planning process, qualitative information in the form of a device being marginally or unfavorably impacted by environmental factors can be very useful in paring the list of potential devices. However, as the mission time draws nearer, quantitative information such as detection and recognition ranges for specific targets becomes more valuable. To effectively present options with respect to the use of these target acquisition devices, a simple tool needs to be available for computing and displaying both types of information. A critical part of this requirement for simplicity is the automated retrieval of environmental input data for the model from a database of values. As part of the Army's Battle Command System (ABCS), the Integrated Meteorological System (IMETS) can provide this database of information. A program has been written to retrieve from the IMETS Gridded Meteorological Database (GMDB) the required current and historical data for a user specified location. The GMDB contains high-resolution (10 km in the horizontal) model data as forecast by a prognostic mesoscale model. This paper presents in more detail the intuitive model and interface that has been developed to allow the user to automatically query the GMDB for the required environmental data and then compute and display the qualitative and/or quantitative information affecting the performance of target acquisition devices. Also reviewed are plans for an Army specific version of the Air Force Target Acquisition Weather Software (TAWS) and ongoing efforts to more fully automate the process of computing and presenting information to the user.

INTRODUCTION

As evidenced by the number and types of papers presented at the annual Ground Target Modeling & Validation Conference, there is a significant amount of basic research being performed related to target signatures, propagation, background characterization, simulation, etc. Perhaps as important to the tactical end user, however, is the design and interface of the resulting target acquisition models. Unless the model presented to this user is intuitive and easy to use and interpret (as well as fast running), the underlying physics associated with the underlying models may not matter since the user may not use the product.

WEATHER TACTICAL DECISION AIDS

Weather tactical decision aids (TDAs) come in two forms: rule-based and physics-based. Rule-based TDAs are constructed using rules that have been collected from field manuals, training centers and schools, and subject matter experts. An example of one such rule would be "usage of TOW2 is not recommended for visibilities less than 3km". Physics-based TDAs employ physics calculations that have their basis in theory or field measurements. Thus a physics-based TDA employs routines and physics that allow it to ascertain the probability of detecting a given target at a given range under existing or predicted weather conditions. This is in contradistinction to the rule-based TDA which would *suggest* that a particular system be used under given weather conditions; no target is required. The trade off between target acquisition (physics-based) and system selection (rule-based) is run time: rule-based TDAs run considerably quicker than physics-based TDAs.

The Army's IWEDA is being adopted as the model for rule-based weather impact decision aids for all the services. A rule-based decision aid provides a general framework based on lists of "if-then-else" rules and pre-established critical weather thresholds for moderate or severe impacts. The Air Force, Navy

and Army, with concurrence from the Marine Corps, are now collecting weather impacts using this common format, and the current database of hundreds of rules are expected to expand to several thousands of rules. The Army IWEDA is designed for the Army Common Hardware/software and DII/COE as part of IMETS, the C4I weather system currently fielded and being improved for the Army First Digitized Division.

Physics based tactical decision aids, as distinct from rule-based decision aids, perform detailed performance calculations for specific systems. Tri-service models for electro-optic propagation, such as the Target Acquisition Weather Software (TAWS), are being linked to rule-based TDA's to provide more detailed effects and quantitative information.

TAWS

The Tri-Service TAWS and its predecessor, the Air Force's EOTDA (the Army's TDA, TARGAC, is a variant of EOTDA), are software models that predict the performance of weapon systems and direct view optics based on environmental and tactical information. Performance is expressed primarily in terms of maximum detection, recognition or lock-on range. The EOTDA and TARGAC supported systems in three regions of the spectrum: visible/TV $(0.4 - 0.9 \ \mu m)$, Laser $(1.06 \ \mu m)$ and far IR $(8.0 - 12.0 \ \mu m)$. TAWS extends this to include the mid-IR $(3.0 - 5.0 \ \mu m)$.

TAWS consists of three essential parts: an inherent target contrast model, an atmospheric effects model, and a system performance model. These basic models are required for any type of device treated by TAWS; however, the nature of the models that perform a given function may vary considerably depending on the system of interest. For example, the inherent contrast for visual and TV devices depends on the relative reflectances of the target and background. Thus atmospheric scattering and solar illumination [Bangert, 1998] are overriding factors. For thermal imagers, the inherent contrast consists of the difference in temperature between the target and background. Solar loading is the dominant factor thus requiring access to previous and forecast conditions required by the thermal balance model [Johnson, et al, 1995, 1998]. The nature of the calculations of atmospheric propagation is also different in the visual than in the IR. Scattering primarily influences the former, while the latter is dominated by absorption. The Army Model and Simulation Office standards category model Acquire is used for sensor performance [Acquire, 1995].

TAWS will also allow for automated meteorological ingest via the Air Force Weather Information Network; these meteorological variables may be selected or changed through numerous Graphical User Interfaces (GUIs). TAWS version 1 is scheduled for release in December 1999 [Tattelman, 1999].

IWEDA

In an effort to simplify the manner in which environmental impacts on weapon systems are displayed to the user, the Army Research Laboratory has developed the rule based Integrated Weather Effects Decision Aid (IWEDA) [Sauter, 1999]. IWEDA provides current and forecast qualitative impacts on approximately 70 weapon systems (e.g., attack helicopters, fixed wing aircraft, personnel, etc) to both meteorologists and non-meteorologists throughout the ABCS (the Army's tactical command and control system). IWEDA "rules" are fired against the GMDB parameters via an IWEDA preprocessor that then populates several relational database tables. A separate program, DIRECT, that allows the Staff Weather Officer (the IMETS operator) to edit and delete specific rules associated with the weapon system has also been developed [Torres, 1998]. DIRECT allows for tailoring of the ruleset to accommodate a specific exercise (e.g., peacetime vs. wartime values) or user (e.g., beginner pilots may have more strict limitations regarding the operation of an aircraft than an experienced pilot). Users configure systems as to how they will be employed, select the systems or missions they wish to see impacts on, then click a button to have the impacts automatically computed via internal queries to the database tables. Once the queries are completed (typically a few to several seconds) a color-coded (red, amber, green) matrix is displayed conveying the impact of the environment on the systems of interest (figure 1). Drill down capabilities are available via mouse clicks such that the user can query and view various levels of information (e.g., condensed or detailed text impact statements, spatial distribution of the impacts via a map overlay [figure 2], etc) depending on their requirements. For example, by clicking the left mouse button anywhere over the map overlay, the full impact statement on the particular item of interest is retrieved and displayed to the user. Thus, it was a simple extension of this concept to allow the computation and display of quantitative target detection and recognition ranges from within the familiar IWEDA GUI. In practice, target acquisition devices within the rule based IWEDA could be evaluated to assess locations and times that particular sensors could be employed
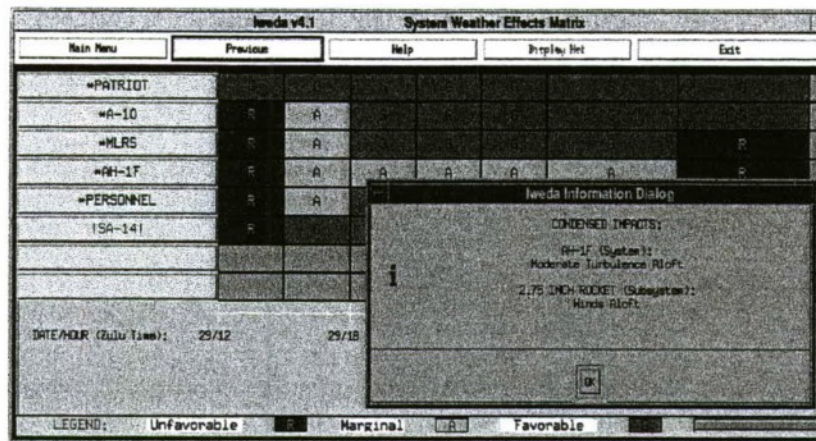
Figure 1. IWEDA Weather Effects Matrix (WEM) with Condensed Impacts
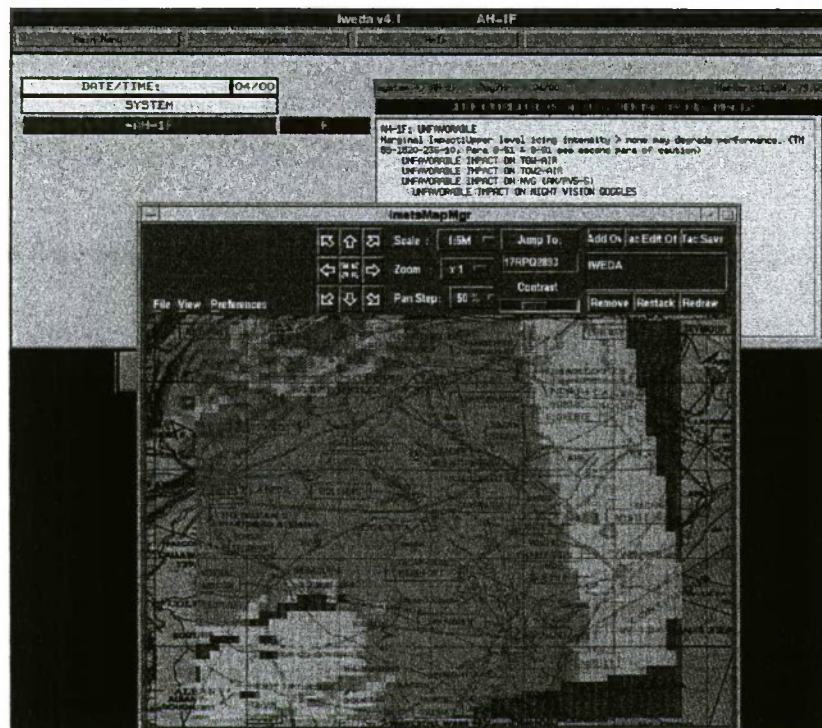
Figure 2. IWEDA Map Overlay for AH-1F (Cobra) Helicopter with Full Impact Statement

with little environmental degradation. For users requiring specific range information, IWEDA will be modified to call the physics based algorithms to compute this information for specific locations on the IWEDA map overlay.

SOFTWARE

Automated Environmental Data Retrieval

Although the target acquisition software has not yet been integrated within IWEDA as a single application, the target acquisition program resides as a standalone application on a Sun Ultra workstation within the development environment at ARL. This application requires a significant amount of environmental data, both for the present and (depending on the sensor) past. Expecting the user to input all of this information when there is gridded meteorological data available is unrealistic. Thus a

front-end module was developed to automatically query, retrieve, and chronologically order the required data. Upon user entry of a valid geographic location and date/time field this information is passed to the data retrieval module. The GMDB typically has data at 10 km horizontal resolution and to a height of approximately 20,000 ft above ground level. The model used to populate the GMDB is the ARL developed Battlescale Forecast Model (BFM) [Henmi and Dumais, 1998] and is typically run twice a day with the forecast time extending to 24 hours. BFM is a hydrostatic model providing for the forecast of numerous meteorological parameters and is currently fielded with the Army's IMETS. A post processing application, the Atmospheric Sounding Program (ASP) [Passner, 1999], derives additional environmental parameters (e.g., turbulence, icing intensity, visibility) from the basic BFM parameters and writes them to the GMDB. The amount of time required to execute the BFM and ASP (to include writing to the GMDB) is typically less than an hour on a Sun Ultra. Archiving of two GMDB data sets (the current and prior run) will allow for gridded data that extends a minimum of approximately 16 hours and a maximum of 28 hours in the past to be available (receipt of model initialization data from the Air Force Weather Agency and model run times add several hours to the time that gridded data is first available). Depending on the parameter being retrieved, the value is obtained either via a weighted interpolation between surrounding grid points (e.g., wind speed) or by assigning the value of the nearest neighboring grid point (e.g., precipitation type). Once successfully retrieved from the GMDB it is passed to the target acquisition routine where the detection and recognition ranges are computed and then graphically displayed to the user.

Interim Target Acquisition Module

Because TAWS version 1 is not currently available, the Army's version of EOTDA, TARGAC, is being utilized to determine target acquisition ranges within IWEDA. The basic differences between TARGAC and EOTDA are the inclusion of Army sensors and targets and the ability to examine surface to air scenarios. As a stand-alone program TARGAC is a part of, and available through the Electro-Optical Systems Atmospheric Effects Library (EOSAEL) [Shirkey, et al, 1987]. TARGAC has been successfully coupled with the GMDB thereby precluding the necessity for inputting past, present and future meteorological data. Coupling with the GMDB also allows target acquisition to be accomplished in near real-time under forecast weather conditions.

Migration and Integration of TAWS-A with IWEDA

Since TAWS is a Tri-Service program, version 1 and all subsequent versions will include Army targets and sensors. The Night Vision and Electronic Sensor Directorate's range performance model for target acquisition systems, Acquire, will be added in version 3 of TAWS. A new atmospheric scattering routine is also being added to TAWS to allow examination of all possible geometrical scenarios (e.g. air-to-ground, ground-to-air, ground-to-ground, and air-to-air). However, in accordance with the Army's IMETS philosophy of keeping input as simple as possible a variant of TAWS will be used in IMETS. This variant will be called TAWS-Army (TAWS-A) and will be functionally equivalent to TAWS. The salient differences will be in allowed user input: TAWS has numerous GUIs to allow automated or manual input and manipulation of meteorological parameters. Since TAWS-A will have automatic ingest of weather conditions through the GMDB there will only be one GUI (figure 3) which will be used to initially select sensor, target

Figure 3. TAWS-A Input GUI

and background information. Currently some additional information is required about the background aerosol. Since the purpose of TAWS-A is to aid the commander make rapid decisions concerning sensor and/or platform selection, the graphical output from TAWS-A will also be different from TAWS. For sensor selection TAWS-A output will be in the form of bar charts (figure 4) or weapon fans (figure 5).
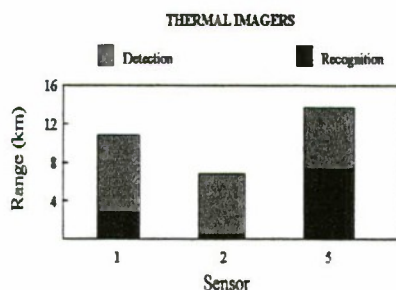
THERMAL IMAGERS

Figure 4. TAWS-A Example Sensor Selection

Efforts are currently underway to integrate TARGAC and, when released, TAWS-A, with IWEDA. As outlined in the introduction, this would allow the user to click on any location on the IWEDA map overlay (figure 2) to have the detection and recognition ranges computed for a specific target/sensor pairing. The default sensor would be for the device currently

Unobstructed line-of-sight
Weather impacted line-of-sight
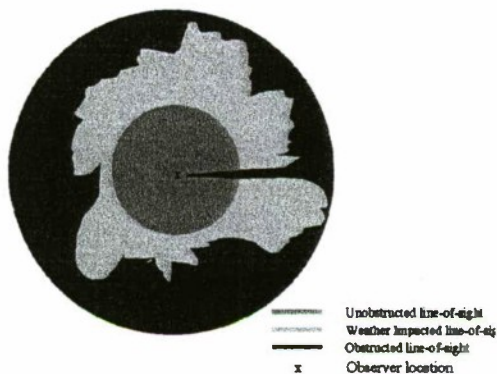Obstructed line-of-sight
x   Observer location

Figure 5. TAWS-A Example Weapons Fan

being evaluated by IWEDA. It is anticipated that upon the user click on the map overlay, a simple GUI would pop up prompting the user for the specific sensor and target. Environmental data required for the computation would be automatically retrieved from the GMDB for the specified point. Allowing the entry of a different device will allow alternate sensors to be evaluated in the event the current sensor is inadequate for the acquisition. In separate windows the detection and recognition ranges would be displayed via bar charts (e.g. figure 4) and the weapons fan would be overlaid on top of the map display (e.g. figure 5). This feature will color code the detection and recognition ranges for numerous azimuths either over a full circle or some smaller arc as specified by the user. Digital terrain elevation data will also be read such that masking due to intervening terrain can be incorporated into the weapon fan display. Eventual incorporation of threat sensors will allow comparison of friendly and threat systems for

aiding in the determination of when and where one force may have the advantage over another.

SUMMARY

Providing a tactical commander with an easy to use and interpret application leveraging the latest target and sensor modeling technology will allow for near real time evaluation of sensor employment options. Automating the environmental parameter retrieval using a prognostic data set further enhances the application by allowing for realistic planning based on evolving weather. Although the current gridded data set currently only extends to 24 hours, a 96 hour capability is anticipated in FY00. Allowing for either standalone running of TAWS or via an integrated environment with IWEDA provides additional flexibility. Leveraging ARL's experience and role in the ABCS, this software will be transitioned to the fielded Army incrementally over the next several years providing an advanced capability where virtually none exists today.

REFERENCES

[1] Bangert, J., SLAC, 1998, http://aa.usno.navy.mil/AA/DoD/software/index.html

[2] Johnson, K. R., Wood B.S., Rynes, P.L., Yee, B.K., Burroughs, F.C., and Byrd, T., 1995, "A Methodology for Rapid Calculation of Computational Thermal Models," *SAE International Congress & Exposition*, Detroit, MI, [http://www.thermoanalytics.com/].

[3] Johnson, K. R., Curran, A., Less, D., Levanen, D., Marttila, E., Gonda, T., and Jones, J., 1998, "MuSES: A New Heat and Signature Management Design Tool for Virtual Prototyping," *Proceedings of the Ninth Annual Ground Target Modeling & Validation Conference*, Houghton, MI, [http://www.thermoanalytics.com/].

[4] "ACQUIRE Range Performance Model for Target Acquisition Systems," 1995, Version 1 User's Guide, *U.S. Army CECOM Night Vision and Electronic Sensors Directorate Report*, Ft. Belvoir, VA.

[5] P. Tattelman, 1999, private communication [tattelman@plh.af.mil].

[6] Sauter, David, et al, 1999, "The Integrated Weather Effects Decision Aid: A Common Software Tool to Assist in Command and Control Decision Making", *Proceedings of the 1999 Command &*

Control Research & Technology Symposium, Newport, RI.

[7] Torres, M.A., 1997, "DIRECT: Dynamic IWEDA Rule Editor and Database Control Tool", *Proceedings of the 1997 Battlespace Atmospherics Conference*, San Diego, CA.

[8] Henmi, T., and Dumais, R. Jr., 1998, "Description of the Battlescale Forecast Model", *Army Research Laboratory Technical Report*, ARL-TR-1032, White Sands Missile Range, NM.

[9] Passner, J., 1999, "The Atmospheric Sounding Program: An Analysis and Forecasting Tool for Weather Hazards on the Battlefield," *Army Research Laboratory Technical Report*, ARL-TR-1883, White Sands Missile Range, NM.

[10] Shirkey R. C., L. D. Duncan and F. E. Niles, 1987, "The Electro-Optical Systems Atmospheric Effects Library, Executive Summary", *Atmospheric Sciences Laboratory Technical Report*, ASL-TR-0221-1, White Sands Missile Range, NM.

RadBase™: A Java-Based Radar Database Generation Toolkit for the Proceedings of the Ground Target Modeling and Validation Conference

Chris Blasband, Jim Jafolla

Surface Optics Corporation

San Diego, CA 92127 USA

ABSTRACT

Predicting the Radar Cross Section (RCS) of complex targets is a complicated problem in electromagnetic scattering theory. Radar scattering theory has been studied extensively for many decades. There have been many software packages developed for calculating the RCS of complex targets, however, those that have been fully validated, and are available to the radar scattering community, require the purchase of a very expensive, high-powered Unix platform.

Surface Optics Corporation has developed a state-of-the-art, commercial off-the-shelf (COTS) software package that disposes of the need for buying an expensive computer. RadBase is a user-friendly software product that generates accurate RCS and Amplitude & Phase data for complex targets and cultural features. It is a Java-based application that executes on PC's running Windows (95,98,NT) and Unix-based platforms. RadBase calculates RCS values using the physical optics approximation to electromagnetic scattering. Additionally, RadBase incorporates physical phenomena such as blocking, double bounce interaction, edge effects, polarization, traveling waves and creeping waves.

In this paper, the capabilities of RadBase 1.0 will be described. A brief description of the scattering phenomena, input and output data is presented. A full RadBase validation study is presented as well as RadBase comparisons with the radar scattering code, Xpatch.

INTRODUCTION

Surface Optics Corporation (SOC) has developed a user-friendly software product for generating accurate Radar Cross Section (RCS) and Amplitude and Phase data for complex targets and cultural features as a function of frequency, polarization, incident angle and azimuth angle. RadBase calculates an object's RCS and Amplitude and Phase data using the Physical Optics approximation to electromagnetic scattering and includes the following effects:

- Blocking
- Double Bounce Interaction
- Edge Diffraction
- Polarization
- Traveling & Creeping Waves

RadBase is a Java-based application and has a user-friendly Graphical User Interface (GUI), with easy to understand input parameters. This allows it to run on Windows 95, Windows 98, Windows NT and UNIX (available Release 1.1, Third Quarter, 1999) operating systems. RadBase currently, supports the following 3-D object model formats:

- STK .mdl Format
- Open Flight (.flt extension)
- ACAD (e.g. Xpatch facet extension)
- Object (.obj extension)

UNIQUE FEATURES

The user does not have to be a radar expert in order to set up and run RadBase. RadBase gives the user:

- Flexibility
- RadBase is a true toolkit
- Users can easily control the program and vary parameters
- Accuracy on a PC
- RadBase generates accurate RCS data for complex targets
- RadBase does not require the purchase of a high-end SGl
- RadBase has been validated against range measurements
- RadBase output compares extremely well with Xpatch
- Speed
- RadBase is very efficient

TECHNICAL DESCRIPTION

The core of RadBase is a software module that calculates the RCS and Amplitude & Phase data for complex targets and cultural features as a function of frequency, polarization and target/observer geometry. This core module uses the physical optics approximation to electromagnetic scattering.

The Chu-Stratton integrals for the total electric and magnetic fields scattered from an object can be very difficult to solve explicitly. High frequency techniques have been developed for solving these integrals. One such technique, Physical Optics (PO) is an approach that is based upon source currents. PO is valid for cases where the incident wavelength is much smaller than the length of the object that is scattering the energy. In PO theory, the geometry of the object becomes very important in calculating the total scattered electric and magnetic fields. PO uses the integral equation representation for the scattered fields. It also uses the high frequency assumption that the scattered field from one point on an object to any other point is negligible compared to the incident field. Therefore the total field at each point on the surface of the object is approximately equal to the incident field at that point. The scattered field is now reduced to a much simpler equation. The surface current density for PO is defined by:

$$\bar{J} = 2(\hat{n} x \bar{H}_{inc})$$

$\bar{J}$ = *surface current density*

$\hat{n}$ = *surface or facet unit normal*

$\bar{H}_{inc}$ = *incident field*

RadBase solves the physical optics integral equations to compute the RCS and Amplitude and Phase for complex objects. Beyond the PO treatment, RadBase has incorporated the following:

- Edge Diffraction Effects
- Blocking
- Multiple Bounces
- Polarization Dependent Scattering (VV, HH) (Cross polarization will be available with Release 1.1, Third Quarter, 1999)
- Traveling and Creeping Waves

RadBase INPUT

A RadBase session begins with the main menu being presented to the user. Figure 1-1 presents a sample of the RadBase main menu.
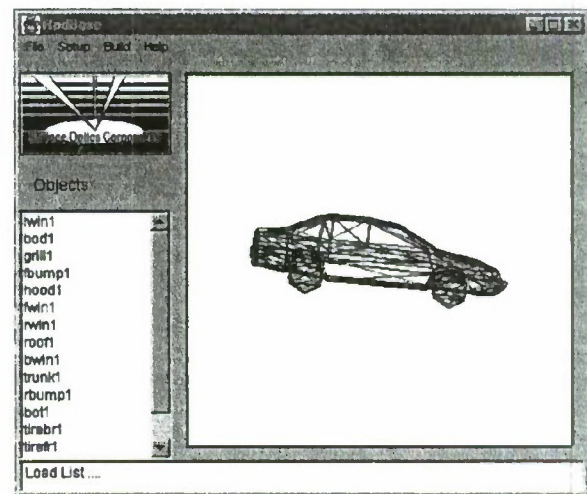

Figure 1-1. RadBase Main GUI

RadBase requires as input:
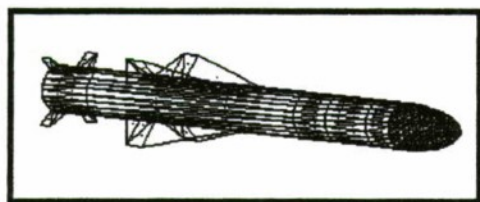
1. 3-D wireframe model of the object

Figure 3-2. Wireframe Model of a Missile

2. User defined radar, geometry and option parameters

User Defined Parameters

Under the "Setup" menu, the user chooses the type of run to be performed:

- Custom

- RadarWorks™ (RadarWorks is a trademark of MultiGen-Paradigm, Inc.)

- STK™ (STK is a trademark of Analytical Graphics, Inc.)

Custom Setup

RadBase requires specific radar, geometry and option parameters in order to develop the output RCS and Amplitude & Phase databases. Figure 1-3 presents the RadBase "Setup Custom" panel.

Figure 1-3. RadBase Custom Setup Window

The user defined input parameters are:

- Target ID
  The Target ID is an integer value placed in the header of the output files. This is useful for simulation systems that contain many targets, allowing the system to uniquely identify each target and assign the proper RCS data.

- Bounces
  This option gives the user the ability to control the number of multiple bounces off of the object that the radar beam can undergo before returning to the receiver. Currently, up to two bounces is allowed, however, Release 2.0 will allow the user to go to *n* bounces.

- Blocking Options (None, Low, Medium, High)
  The speed of RadBase is dependent upon the number of object facets, the number of bounces, and how many blocking computations are desired. Blocking refers to one facet of the vehicle blocking another facet from the radar view. If an object is fairly simple and has no facets that can block another, than it is recommended that the user input "None". If the object is very complex and has many facets that can block other facets, the user should input "High". The remaining two options are for intermediate objects. This is truly a subjective input by the user and for the best "physics" treatment, the user should use "High" if unsure.

- Output Option: Binary or ASCII
  The output database can be either binary or ascii.

- Edge Diffraction (Yes/No)
  Edge effects can have a significant effect on the RCS of a complex object. RadBase includes edge diffraction effects coupled with the physical optics computations. RadBase is delivered with a routine that automatically generates object edges. It uses the information in the vertex and facet file to generate these edges. If edge diffraction is chosen, RadEdge generates edge information that is passed directly into the RadBase edge diffraction routine.

- Maximum Interior Wedge Angle (MIWA) (Degrees)
  The maximum interior wedge angle defines the maximum angle at which two facets form an edge. Angles greater than this angle are not included in
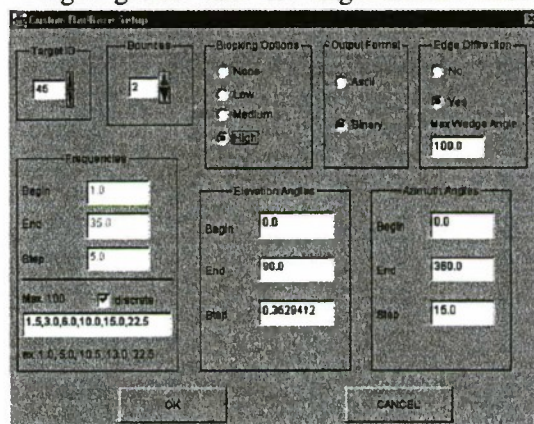
the edge computations. This prevents inclusion of edges in which the interior angle is close to 180°. The default value for this parameter is 100°. It is recommended that the user choose values between 100° and 170°. The limits are:

$$90° < MIWA < 180°$$

- Frequencies (GHz) (1-100 GHz)
  - Uniform Frequencies (Begin, End, Step)

    The user can input a begin, end, and step frequency. The uniform frequency option is useful for generating RCS and Amplitude and Phase data to be used in a SAR processor, or if a known desired frequency spacing is desired. The maximum number of frequencies is 100.

  - Discrete Frequencies
    The user can input up to one hundred (100) discrete frequencies in any given RadBase run.

- Elevation Angles (Begin, End, Step)
  The elevation angles must be input in degrees. The range of angles is from 0° to 360°. 0° - 180° defines the upper hemisphere and 180° - 360° defines the lower hemisphere. The maximum number of elevation angles is 2001.

- Azimuth Angles (Begin, End, Step)
  The azimuth angles must be input in degrees. The range of angles is from 0° to 360°. Rotation is defined as counter-clockwise for an observer sitting above the target. Thus, if the nose of the aircraft points down the positive x-axis, 90° is viewing the left wing as defined by the pilot and 270° is viewing the right wing. The maximum number of azimuth angles is 2001.

RadBase OUTPUT

RadBase outputs two binary or ascii files. The prefix of the input file is used for output. For example, if the file, "barn.flt" is input, the two output files will be named, "barn.rcs" and "barn.aph". The .rcs file contains RCS data for all user-defined frequencies, two linear polarizations and all user defined elevation and azimuth angles. The .aph file contains amplitude and phase data for all user-defined frequencies, two linear polarizations and all user defined elevation and azimuth angles. Figure 1-4 presents an RCS versus elevation angle plot generated by RadBase.
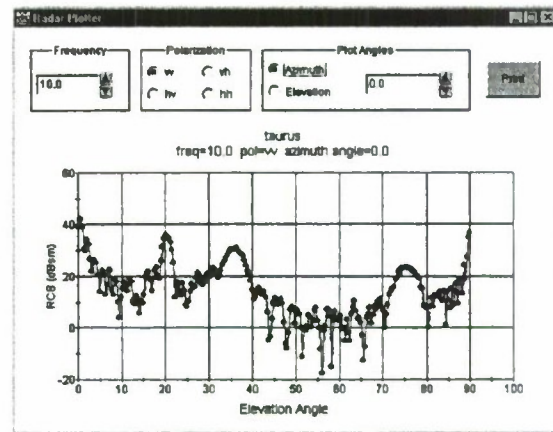
Figure 1-4. RCS vs. Elevation Angle Plot Generated
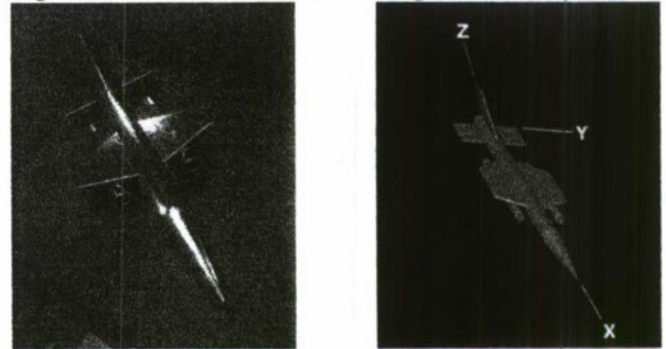
by RadBase

VALIDATION

RadBase has been validated against range measurements of an aircraft. Figure 1-5 presents an image of the aircraft measured in the radar range and Figure 1-6 presents an image of the modeled aircraft.

Figure 1-5. True   **Figure 1-6. Modeled**

Figures 1-7 through 1-10 present plots of the RadBase predicted RCS versus the measured data. It can be seen that RadBase compares extremely well with the measured data. The discrepancies beyond 300° for the 10 GHz and 15 GHz, 30° elevation cases are due to the fact that RadBase does not yet treat cavities or calculate more than two bounces. Release 2.0, scheduled for Fourth Quarter,

5 GHz; VV; Elevation = 42 Degrees
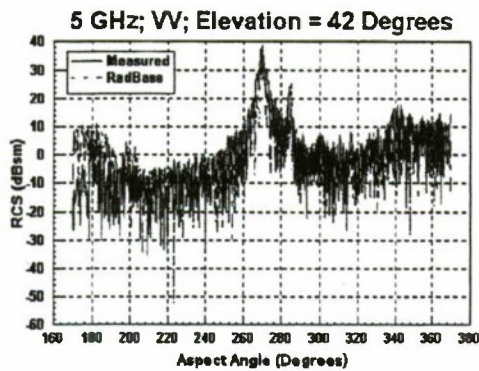
1999 will have a treatment for cavities as well as include a robust model for calculating the RCS for any user-specified number of bounces.

5 GHz; HH; Elevation = 42 Degrees

10 GHz; VV; Elevation = 30 Degrees
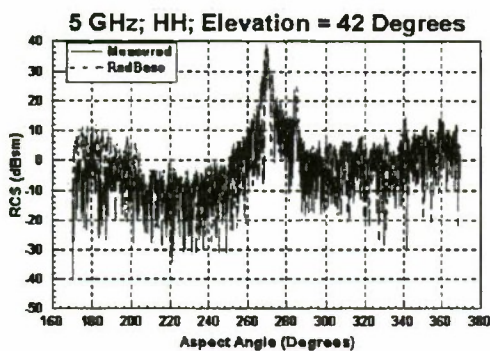
Figure 1-7. Measured vs RadBase Modeled Aircraft 5 GHz; VV; 42° Elevation

Figure 1-8. Measured vs RadBase Modeled Aircraft 5 GHz; HH; 42° Elevation

Figure 1-9. Measured vs RadBase Modeled Aircraft 10 GHz; VV; 30° Elevation

15 GHz; HH; Elevation = 30 Degrees
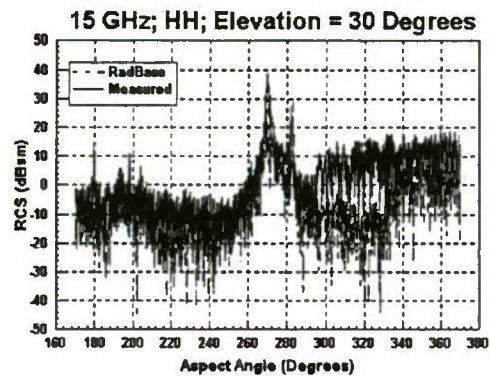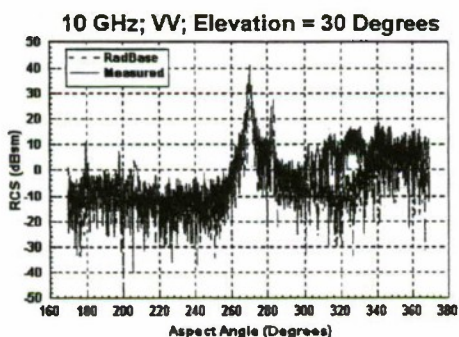
Figure 1-10. Measured vs RadBase Modeled Aircraft 15 GHz; HH; 30° Elevation

Surface Optics is further validating RadBase against range measurements of a T72. Preliminary results show that RadBase is agreeing extremely well with the measured data. The T72 validation report will be available to the public September 1999.

COMPARISON WITH Xpatch

RadBase has been compared to the radar scattering software package, Xpatch. Xpatch is a set of high-frequency radar signature prediction codes that are based on a method called the Shooting & Bouncing Ray (SBR) technique. Xpatch was chosen for this study because it has been well validated and documented.

Figure 1-11 presents a three dimensional image of a Scud Launcher used as input into RadBase and Xpatch. Figures 1-12 and 1-13 present comparisons of RadBase and Xpatch for the Scud Launcher. The figures show that RadBase compares extremely well with Xpatch, and the two produce almost exact results at the major peaks (0°, 76°, 90° and 180°). Figure 1-14 presents a comparison of RadBase and Xpatch for a faceted missile. Again, the two software systems produce very comparable results.

These results are only a small sample from a complete RadBase/Xpatch study performed by Surface Optics.
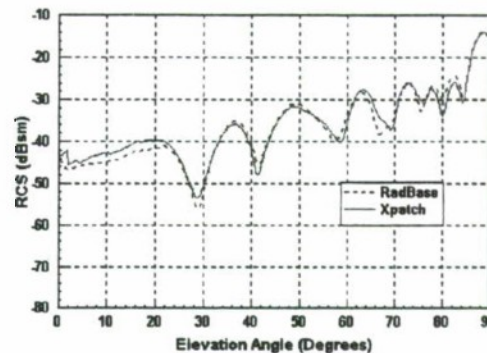
Figure 1-12. RadBase/Xpatch Comparison for a Scud Launcher – 10 GHz; 2=30°; VV

Figure 1-11. Scud Launcher used for RadBase/Xpatch

Comparison

10 GHz; 2=30°; HH

Figure 1-13. RadBase/Xpatch Comparison for a Scud Launcher – 10 GHz; 2=30°; HH

10 GHz; M=0°; VV

Figure 1-14. RadBase/Xpatch Comparison for a Faceted Missile – 10 GHz; M=0°; VV
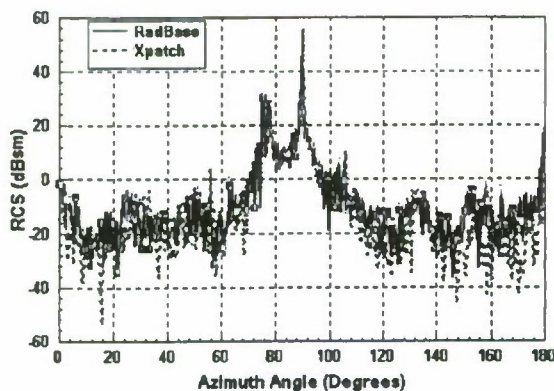
APPLICATIONS

RadBase has applications in many areas of radar research, design and simulation. It is currently being used by radar design engineers, radar experts in the real-time visual simulation/sensor simulation community, and scientists performing Synthetic Aperture Radar (SAR) image interpretation and simulation.

RadBase gives the user the ability to generate radar databases for a variety of applications:

- Radar Simulation
- Target Signature Analysis
- Radar System Analysis
- Radar System Performance
- Radar Design
- SAR Image Interpretation
- Human Factor Studies
- Radar Operator Training
- Input to other COTS Products
  - STK/Radar™
  - RadarWorks™

FUTURE ENHANCEMENTS

RadBase will advance significantly with new releases. Upgrades will include:

- Target Viewer/Material Editor
  - ➤ The user will be able to view a 3-D object, point and click on individual facets or groups of facets and easily make materials assignments (with no typing required)
  - ➤ RadBase will be delivered with a complete set of complex dielectric properties databases
    - ❑ Developed at SOC
    - ❑ Continually enhanced by SOC
- Higher Order Bounces (>2)
- Ability to handle a rough surface ground plane
- Treatment of Cavities

CONCLUSIONS

RadBase is a user-friendly software product for generating accurate Radar Cross Section (RCS) and Amplitude and Phase data for complex targets and cultural features as a function of frequency, polarization, incident angle and azimuth angle. It allows labs to produce radar simulation data on every PC available within their facility. Previously, this was only possible using their handful of very expensive UNIX platforms.

RadBase has been validated against range measurements and been compared extensively to validated radar scattering software packages such as Xpatch. RadBase will continue to advance significantly with new releases.

REFERENCES

[1] Eugene Knott, John Shaeffer, Michael Tuley, "Radar Cross Section", Artech House, Inc. 1985

[2] George Ruck, Donald Barrick, et al, "Radar Cross Section Handbook", Plenum Press 1970

Radar Modeling With Irma

Michael A. Richards, Douglas A. Vechinski, and John Watson
Nichols Research Corporation
1130 Eglin Parkway, Suite A
Shalimar, FL 32579
850-796-3932
richards.mike@nichols.com

Capt. Vince Underwood
Department of the Air Force
Air Force Research Laboratory, AFRL/MNGG
101 West Eglin Blvd
EAFB, FL 32542

Classification: Unclassified

ABSTRACT

Irma is an Air Force sponsored tool for generating synthetic signatures and imagery of complex environments for a wide variety of different sensors and spectral bands. The Irma code has three main signature/sensor components consisting of the passive, ladar, and radar models. These models allow a sensor to view a complex scene consisting of targets, terrain, and atmosphere. There is essentially no limitation on the placement of the sensor or the geometrical complexity of the scene. The passive code accounts for reflected sky/solar/earth radiation, emitted radiation and atmospheric transmittance and path radiance to produce radiance (or apparent temperature) imagery. The ladar code accounts for monostatic surface scattering, atmospheric effects and sensor motion to produce range and cross section imagery. The radar code accounts for spatially coherent and incoherent scattering, atmospheric transmittance and volumetric backscatter and is designed to produce synthetic aperture radar (SAR) and realbeam radar imagery, pulse-doppler maps, and range-resolved cross sections. This paper will focus on the radar model, the associated phenomenology and example signatures, and the recently added pulse-doppler and range-resolved cross section capabilities.

INTRODUCTION

Irma is a mature multi-sensor signature prediction code and analysis tool [1]. The purpose of the Irma radar channel is to generate signatures of target and terrain scenes viewed with a radar sensor [2]. The generated radar signatures typically represent the radar cross section (RCS), or some related quantity, as a function of various (range, doppler, cross-range) coordinates.

The radar model currently has three main types of output signatures: SAR imagery, pulse-doppler maps, and range resolved cross sections. In the SAR mode, the output signature is given as a function of down-range and cross-range coordinates. In the pulse-doppler mode, the return signal power in each range doppler cell is computed. This mode was designed for modeling air-to-air engagements. The code can also output a complex range-resolved return signal where the output is proportional to the square root of the RCS in each range bin. These range profiles are characteristic of air-to-surface real-beam radar systems.

The signature generation process utilized in the Irma code consists of several different steps. First, it is necessary to describe the scene. The scene is described by specifying the location and orientation of all scene features and objects and their scattering properties. Next, the sensor is described by its position, orientation, frequency, resolution, antenna pattern and polarization. The code then determines which objects within the scene

are illuminated by the sensor. In this stage, oversampling on selected objects may be utilized to provide a high-resolution description of the sampled geometry. The radar signature is then determined utilizing a combination of spatially coherent and incoherent techniques. Furthermore, the simulation may utilize a single bounce Z-buffer approximation [3] or, in some cases, the XpatchT [4] module may be used to generate signatures of objects containing multiple bounce components. Also, the model accounts for the sensor blur and speckle effects as well as atmospheric attenuation and backscatter.

The Irma model is written in FORTRAN and C and is supported on both the SUN (Solaris 2.5 or later) and SGI (Irix 6.2 or later) platforms. In addition to these platforms, the radar module, as outlined in this paper, is in the process of being ported to the PC/Windows NT platform. The Irma software is available free of charge to US government-approved personnel.

SAR MODELING ENHANCEMENTS

A SAR image is a graphical display of the radar cross section (or some similar quantity) as a function of down-range and cross-range coordinates. Objects that appear bright in the scene are scattering a relatively large portion of the incident electromagnetic energy back toward the sensor. Surfaces that are highly reflective will appear bright when oriented perpendicularly to the sensor line of sight. Smooth objects that lie flat in the scene, such as roads and water regions, will generally appear dark since they scatter the incident energy in the specular direction away from the sensor. The down range resolution is controlled by the bandwidth (effective pulse width) of the radar. The cross-range resolution for real beam imagery is dependent on the beamwidth and range. For synthetic aperture systems, the cross range resolution is dependent on the mode of operation and effectively how long the sensor views the scene. With Irma, a variety of resolution functions (sinc, Hann, Gaussian, etc.) are supported. These functions may be used to control the resolution and sidelobe effects. Shown in Figure 1 is a SAR image of a rural scene containing fields, a road, trees, and several tanks.
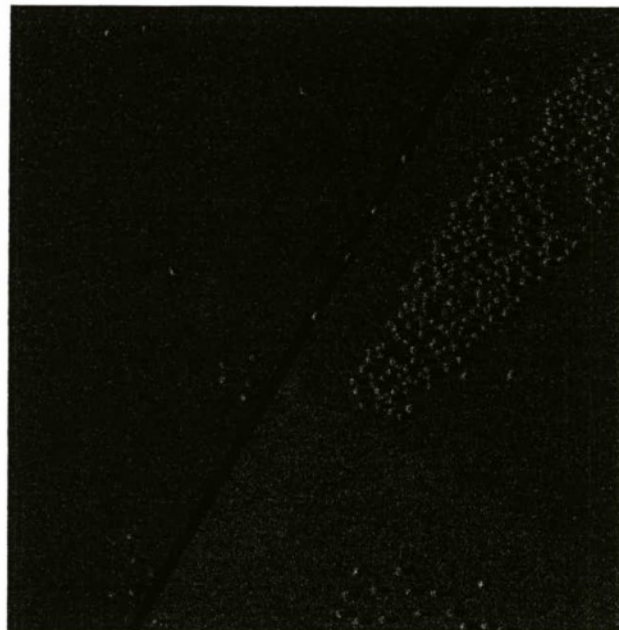
Figure 1. SAR Image of rural scene containing a road, fields, trees, and several tanks.

In addition to the standard imagery that is produced by Irma, output SAR imagery for a set of amplitude monopulse antenna patterns can also be generated. When in this mode, the code utilizes a user-supplied sum antenna pattern, a delta-azimuth pattern, and a delta-elevation pattern. An example of these patterns is shown in Figure 2. The resulting images, such as those shown in Figure 3, can be utilized to determine the angular position of a target within the scene.
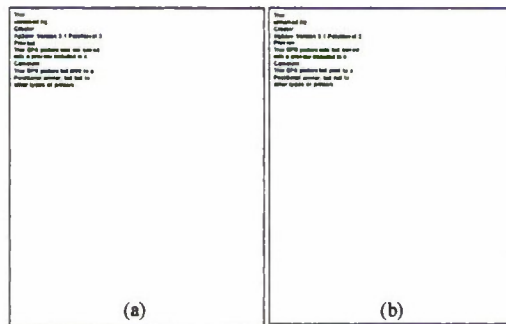
Figure 2. Monopulse (a) sum and (b) delta-azimuth pattern.
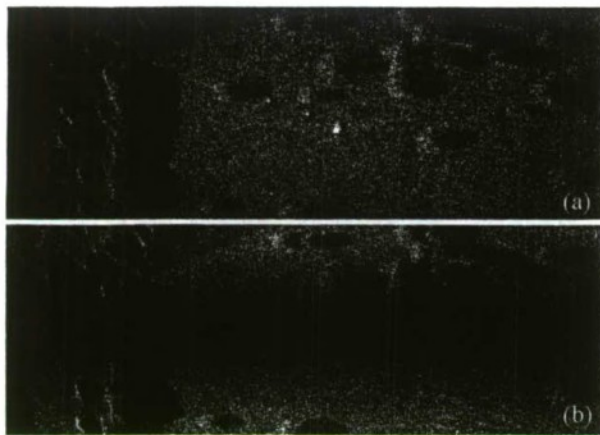
Figure 3. Example monopulse SAR images (a) sum pattern, (b) delta-azimuth pattern.

An important upgrade to the radar channel is the ability to make use of external SAR interpolation tables. Prior to this enhancement, the radar code had the ability to incorporate SAR signatures that were obtained by external means such as Xpatch or possibly even from measurements. However, this was only valid for a given target aspect. If a sequence of images was to be generated which involved a changing target or sensor position, it was necessary to generate the target chips one at a time and manually change the external SAR images each time.

For illustration purposes, a set of SAR images of a tank, shown in Figure 4, was generated using the Xpatch code for several different azimuth angles. The elevation was held constant at 24°. A subset of the data is displayed on a linear scale in Figure 5 with the image minimum and maximum mapped to black and white, respectively.

Figure 4. T72 facet model.
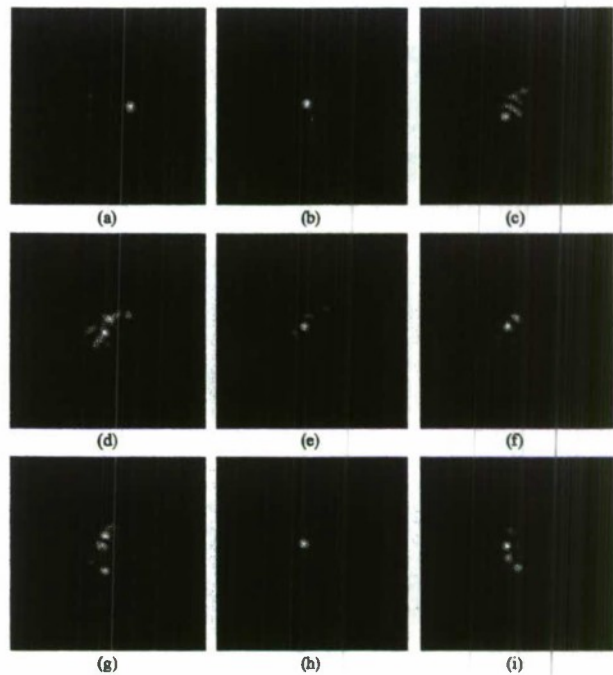
Figure 5. Synthetic Aperture Radar Images of a T72 tank target as viewed from a 24° elevation. The Ka Band sensor aspects are (a) 0° (b) 10° (c) 20° (d) 30° (e) 40° (f) 50° (g) 60° (h) 70° (i) 80°.

The radar code now has the capability to read in a database of filenames which correspond to a number of externally generated SAR images for a number of different aspect angles. Then during the signature calculation, the code determines the appropriate aspect angles of the sensor in the base coordinate system of each target. This location is then used to find the four surrounding images in the database. These images are utilized to obtain a new external image which is then used for the current target. The user also has the capability to scale the image to represent some desired cross section as an aid in performing parametric studies.

The use of the interpolation table is faster and much less prone to error than manually updating the external file. Furthermore, it is much more practical than attempting to call Xpatch each time to generate an image which may take several hours to calculate. However, it is necessary to generate the database of images offline, which will be a time consuming task. An example of the process is shown in Figure 6. Here, the signature is desired at an aspect of $\theta = 73°$, $\phi = -36°$. The table contains signatures for every 5° in elevation and every 10° in azimuth. The surrounding pre-calculated images at $\theta = 70°$ and 75° ,

ϕ = -40° and -30° are then used to obtain an image at the desired aspect.
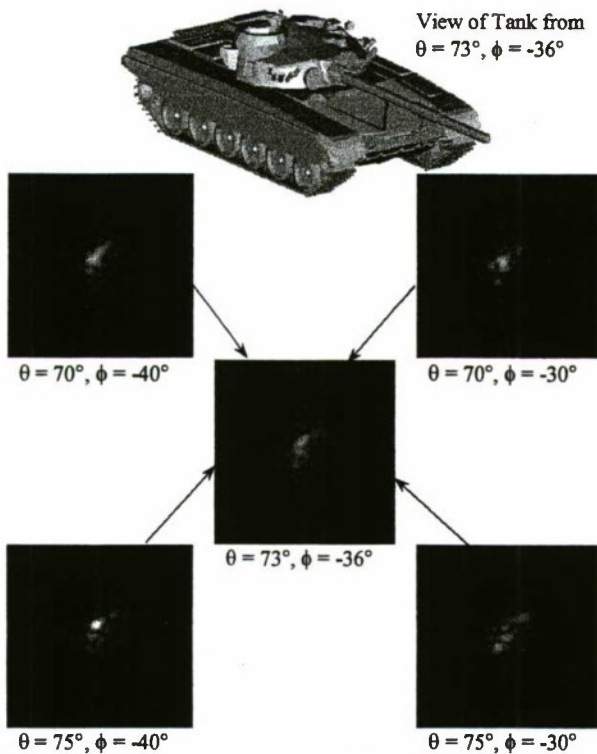
Figure 6. Example of external SAR image interpolation.

For demonstration purposes, an example of incorporating an XpatchT result into Irma is shown in Figure 7. The frequency was set to 35 GHz, the sensor is located at θ = 70° and ϕ = 90°, and only one bounce is used. The external SAR image was obtained from an XpatchT run. This was then incorporated as the top tank in the final image. For comparison, the bottom tank in the image was calculated by Irma with target oversampling.
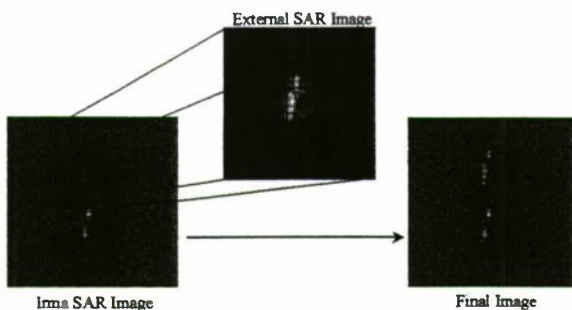
Figure 7. Demonstration of the process of incorporating an external SAR image into Irma. The external Ka Band SAR image of a tank was computed by XpatchT. The lower tank image was computed by Irma.

PULSE-DOPPLER AND RANGE PROFILE ENHANCEMENTS

In addition to the SAR modeling capabilities of Irma, the radar channel also has both pulse-doppler and range-profile sensor capabilities. The primary purpose of the pulse-doppler mode is for modeling air-to-air engagements. As in the SAR mode, the targets and clutter features are described by specifying their location, orientation, and scattering properties. In addition to these properties, the velocity (translational and rotational) of all moving targets is specified. The sensor is specified by its position, orientation, frequency, resolution, antenna pattern, noise characteristics, and translational velocity. The imaging scenario is illustrated in Figure 8.
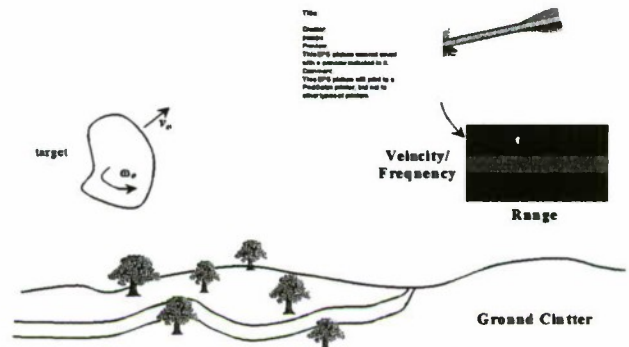
Figure 8. Range-doppler modeling scenario.

The range-doppler maps are created by first rendering the scene. The range to each pixel is determined along with its relative radial velocity with respect to the sensor. These two quantities determine where the signature will get mapped in the radar image. The received power is then determined using the radar range equation and accounts for the range to the target, the pattern, gain, and polarization of the sensor, atmospheric and other loss factors. The primary output of this sensor mode is a range-doppler image or map where the values in this map represent the received signal power. For illustration purposes, a pulse-doppler map is shown in Figure 9 and contains target, clutter, and noise effects.
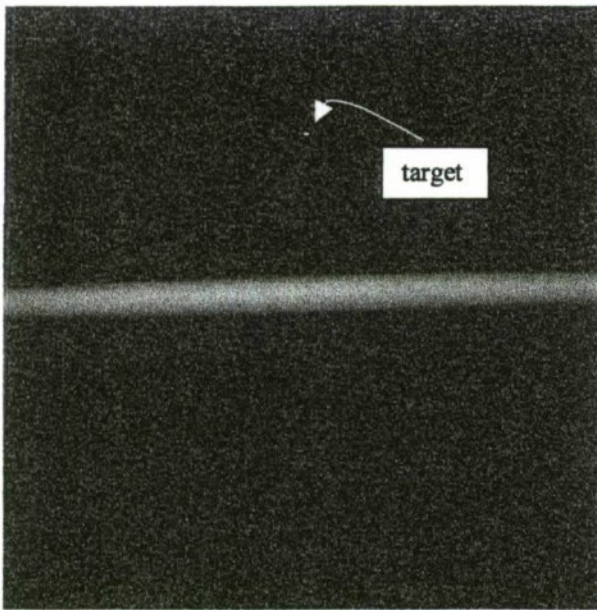
Figure 9. Pulse-doppler image of a scene containing both a target and clutter.

The intended purpose of the range-profile capability in Irma is to model air-to-ground radar guidance and fuzing. In a fuzing scenario, a sensor guides a weapon to the target, and a proximity fuze detonates an explosive warhead when the weapon reaches the planned position, range, or altitude with respect to the target. As the munition closes on the target, a range is reached at which far-field scattering conditions no longer apply. In the near-field region, the phase of the incident field is no longer planar, the range is not constant over the target, and the antenna pattern varies over the target scene. As the sensor transitions from the far-field region to the near-field region, the peak RCS is typically reduced, the nulls are not as deep and the sidelobes are higher.

In performing the fuze calculations, as with the pulse-doppler computations, the sensor is modeled with an arbitrary antenna pattern. This pattern is described by the magnitude and phase of the E_θ and E_ϕ components of the electric field radiation pattern as a function of θ and ϕ in the antenna's local coordinate system.

For illustration purposes, a synthetic range profile of a simple target on a grass background is shown in Figure 10. In this scenario, the Ka-band sensor is positioned 50 m from the target at an elevation of 30°.

Figure 10. Synthetic range profile produced by Irma.

SUMMARY

Irma is a mature multi-sensor signature prediction code. This paper has outlined the status of the Irma radar module with its recent enhancements and presented example SAR imagery, pulse-doppler maps and range profile signatures. These results serve to demonstrate the wide range of scenarios for which Irma can be utilized.

ACKNOWLEDGEMENTS

REFERENCES

1. M. Owens, et al., "Irma 5.0 Multi-Sensor Signature Prediction Model," SPIE Aerosense Conference, April, 1999.

2. M. A. Richards, D. A. Vechinski, and J. Watson, "Radar Image Simulation Via the Irma Code," Ground Target Modeling and Validation Conference, August, 1996.

3. M. A. Richards, "Partial Survey of Codes for High Frequency Scattering from Facet Models of Radar Targets," *Applied Computational Electromagnetics Society Journal*, Vol. 9, No. 2, 1994.

4. S. W. Lee, et al., "User's Manual for Xpatch," Technical Report, DEMACO, June 1997.

Software to Convert Mechanical Desktop Solid Models into Facet Files for Input to Xpatch

Roger Evans, John G. Bennett and Jack Jones
U.S. Army Tank-automotive and Armaments Command
Warren, MI 48317-9000

ABSTRACT

Mechanical Desktop, a solid modeling program from AutoDesk, offers engineers a powerful environment for designing ground combat vehicles. To predict the radar signatures of concept vehicles created in Mechanical Desktop, the geometry and material properties of the vehicles must be fed into Xpatch. Mechanical Desktop, however, lacks an output file format that is directly useable by Xpatch. In this paper, we discuss the problems associated with converting Mechanical Desktop files into a suitable facet format, and we present a procedure for carrying out the conversion. Moreover, we describe new software required to perform a key step in this conversion. The conversion of an example of Mechanical Desktop solid models illustrates the procedure and the new software.

INTRODUCTION

At the U.S. Army Tank-automotive and Armaments Command (TACOM), engineers in our team use the commercial computer aided design (CAD) program Mechanical Desktop to design ground vehicles as 3 dimensional solids. Mechanical Desktop, a product of AutoDesk, offers the vehicle designer a host of features to facilitate the design process and to produce useful output. For example, Mechanical Desktop can transform a solid model into a fully dimensioned AutoCAD drawing for use in the machine shop.

If we could use this geometric description to predict a vehicle's radar signature with Xpatch, then we could avoid the time and labor required to duplicate manually the description of the vehicle. Moreover, using the same geometry description would ensure precise configuration control. Mechanical Desktop, however, does not produce a geometry file that can be fed directly into Xpatch to predict radar signatures.

In this paper, we present a procedure to solve this problem by using a combination of commercial and custom software to convert Mechanical Desktop output files into facet files for input to Xpatch.

CONVERSION PROCESS

Figure 1 illustrates the overall relationship of the programs involved in the conversion process. First, Mechanical Desktop creates the vehicle geometry and exports it as an IGES file. The commercial program Rhino imports the IGES file and exports it as an OBJ file. And, finally, the custom program OBJ to Facet Converter creates the facet file suitable for input to Xpatch.

Rhino, a product of Robert McNeel and Associates (www.rhino3d.com), specializes in the creation of 3 dimensional objects represented as non-uniform rational B-splines (NURBS). But we chose Rhino for this conversion process because it supports a wide range of input and output formats.

Roger Evans wrote the OBJ to Facet Converter, the other key program in the conversion process, in C and Borland C++Builder to run under Windows. The program can be either command line driven or launched from a Windows graphical user interface. Computer memory sets the only limit on the size of the files that can be converted.

To handle the assignment of material codes, the user must split a vehicle into separate files for each material. Figure 2 illustrates the process of converting the separate files and recombining them into an Xpatch vehicle file.

DTANK: AN EXAMPLE OF THE CONVERSION PROCESS

The conversion of DTANK, an example geometry, highlights the features of the conversion process. Figure 3 shows DTANK in Mechanical Desktop, where Dr. David Hansen created the geometry for this paper. The name derives from David's Tank. The colors of DTANK represent different materials, each of which will be converted separately. From Mechanical Desktop, DTANK is exported as IGES files.

In Figure 4, Rhino displays DTANK after the IGES files have been imported. At this point, DTANK is still represented as solid objects. To export the DTANK files as OBJ files, the user can control meshing with a simple slider, Figure 5, or the user can choose detailed meshing parameters, Figure 6. The meshing parameters can have major effects on the number and shape of the facets. For example, the result of selecting different values for the Max. Angle meshing parameter is displayed in Figure 7.

The next step is to run OBJ to Facet Converter, Figure 8. Note that the user has the option to assign an Xpatch material number to the entire file. A dialog box, Figure 9, reports on the results of the conversion.

The separate material facet files must now be combined into a single vehicle file. The files can be combined either in one step by the Combiner program written for this paper, or the files can be combined in several steps, two files at time, by the file combining program packaged with Xpatch. The result of the combining is a single vehicle facet file with components labeled with the proper Xpatch material number, Figure 10.

The file is now ready for input to Xpatch. Figure 11 shows a synthetic aperture radar image of DTANK generated by Xpatch from the converted file.

CONCLUSIONS

The process described here can convert geometries created in Mechanical Desktop into Xpatch compatible facet files. Presumably, this same process could be applied to IGES geometry files created by other CAD programs.

Future improvements might include removing internal facets. Additionally, better techniques are needed to deal with small parts such as bolt heads.

ACKNOWLEDGEMENT
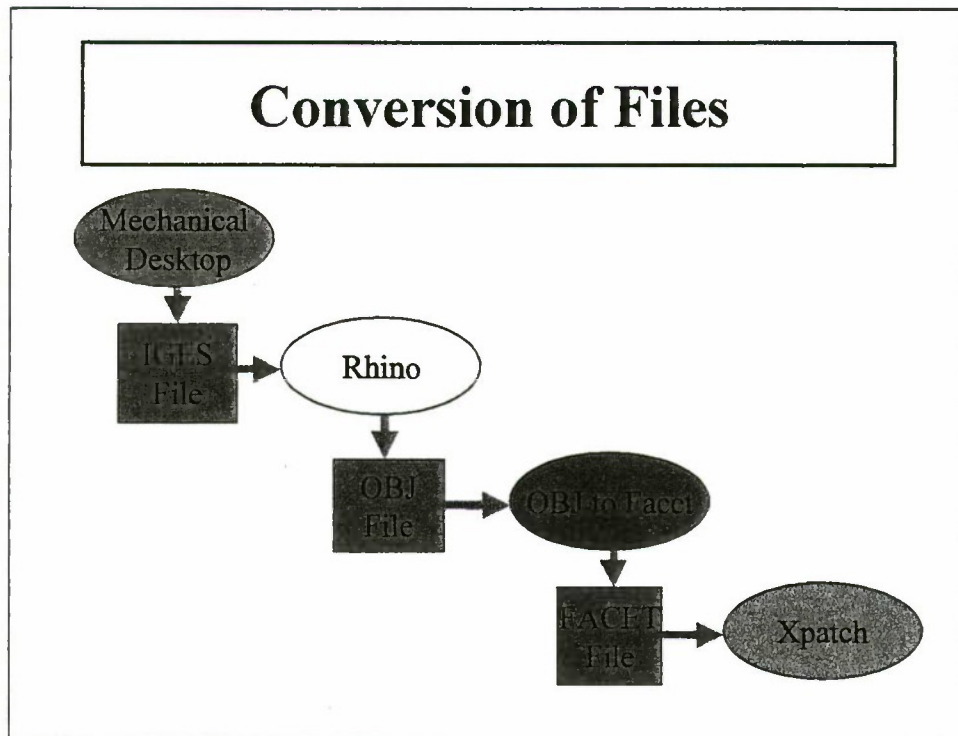
Figure 1. The process of converting a file from Mechanical Desktop to Xpatch.
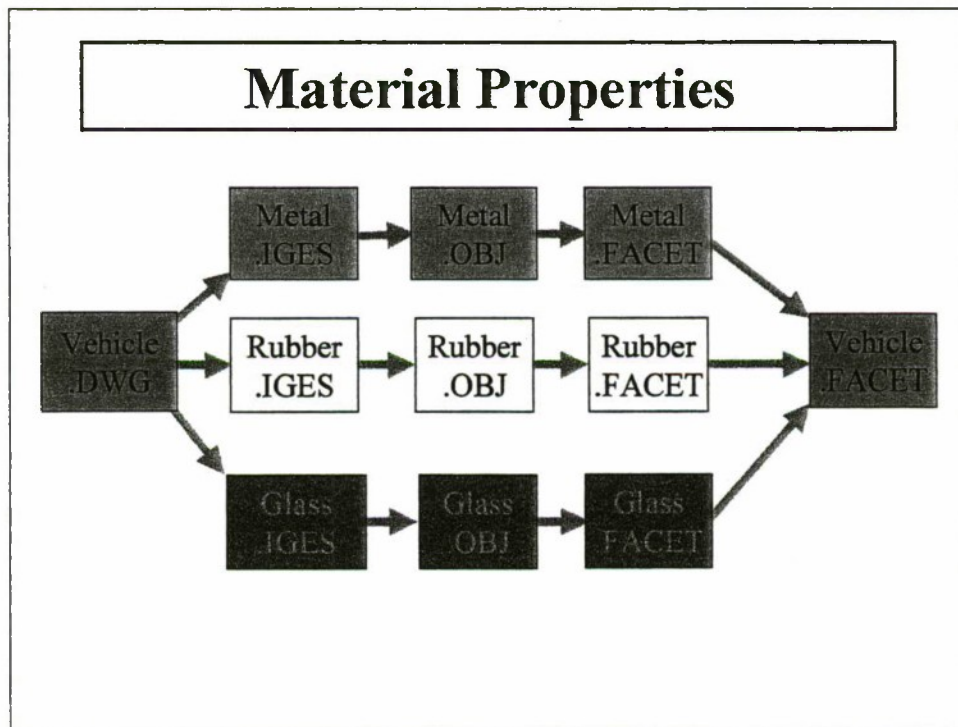
Figure 2. Parallel conversion of components of different materials.

Mechanical Desktop
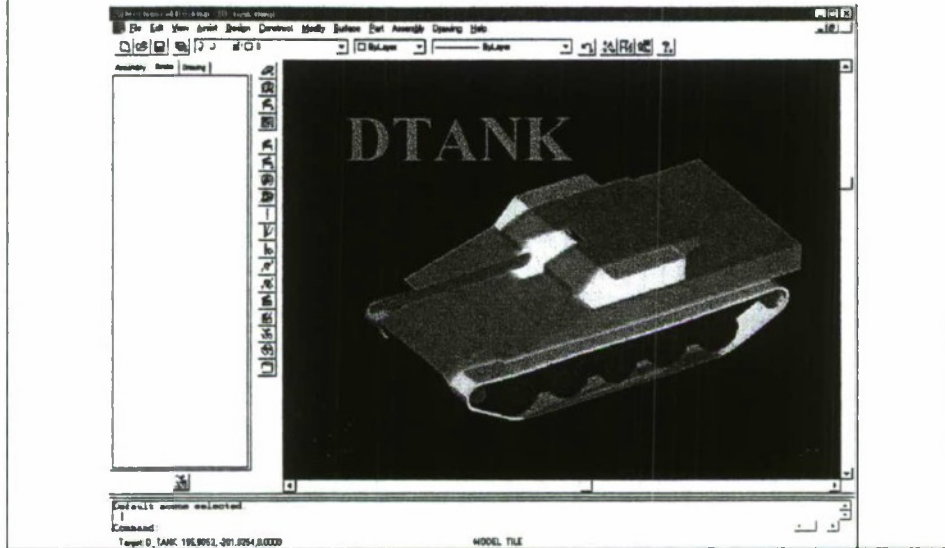
Figure 3. DTANK in Mechanical Desktop.

IGES File Imported into Rhino
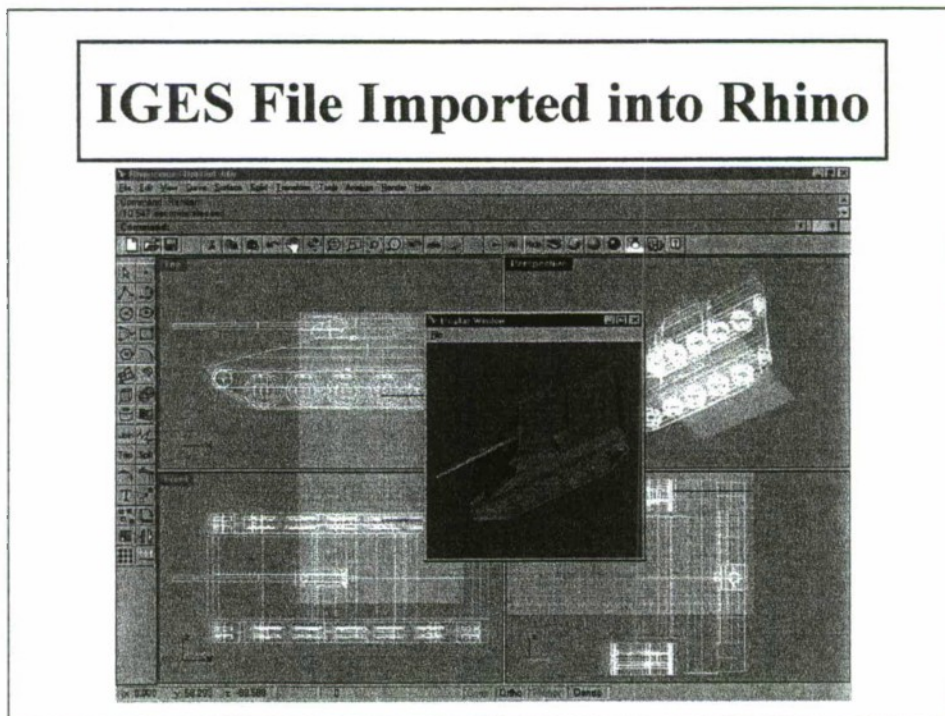
Figure 4. DTANK imported into Rhino.
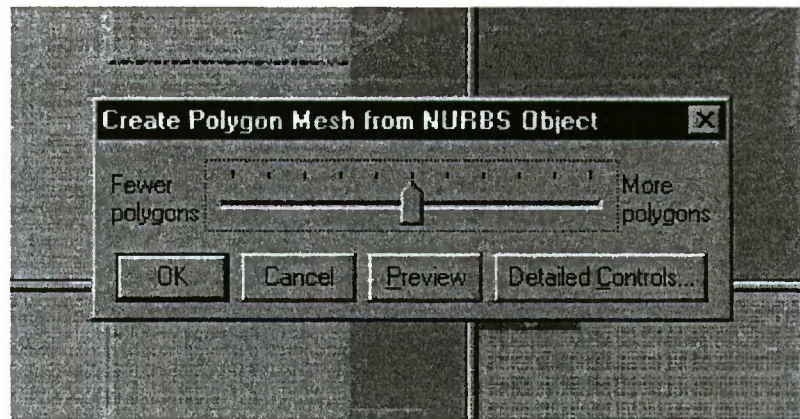
Figure 5. Simple control of meshing parameters for OBJ file.

Figure 6. Detailed control of meshing parameters for OBJ file.
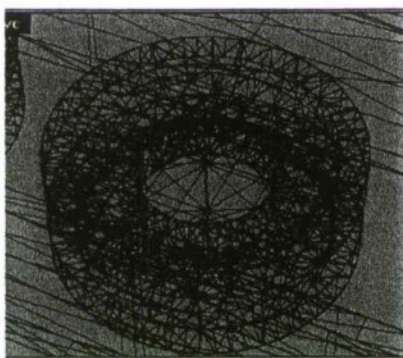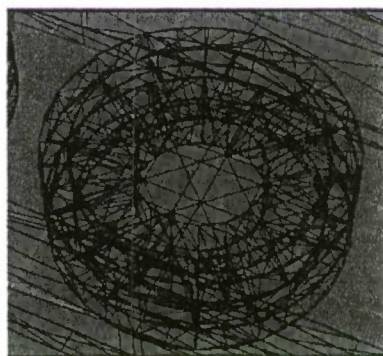
Figure 7. The result of 2 values of the Max. Angle meshing parameter.
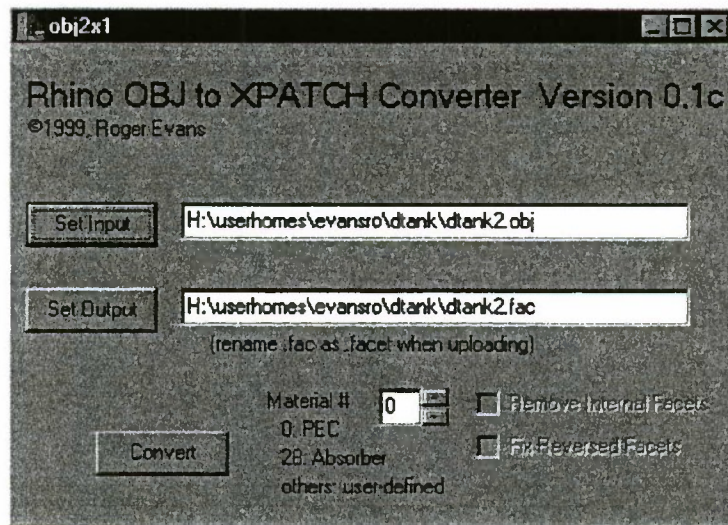
Figure 8. Dialog box to run the OBJ to Facet Converter.
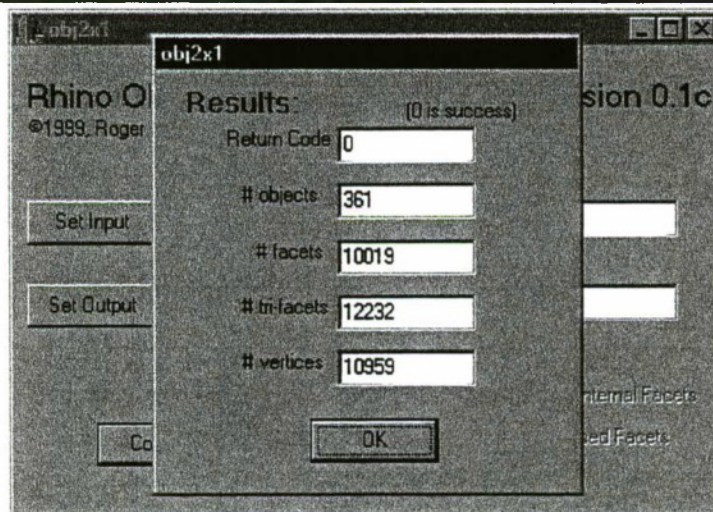
Figure 9. Report on the results of an OBJ to Facet conversion.

Figure 10. The converted vehicle facet file displayed in Xpatch's viewer.

Figure 11. Output from Xpatch generated from the converted vehicle file.

The Benefits of Additional Primitives used by Radar Target Signature (RTS)
for the Proceedings of the
Ground Target Modeling and Validation Conference

Daniel Halstead
ANGLE Incorporated
Springfield, VA 22150 USA

ABSTRACT

In analysis efforts this author has participated in over the last few years, over 60 percent of total effort was spent on the production of a reasonable and accurate mesh for analysis. A comparison (production times and electro-magnetics) of fast automeshing triangles (higher element counts and longer analysis run times) with reduced element count models using additional primitives (longer mesh preparation time and reduced analysis run times) would be of interest. The target will be a truncated pyramid with quarter wavelength cracks, fastener bumps and weld beads. The approach involves first, an automesh using only triangular elements, and second is a semi-automated mesh using quadrangles, cylinders, cones, elliptical plates, ellipsoids and paraboloids as well. The analysis, accomplished using NRL's Radar Target Signature (RTS), is comprised of full azimuthal runs at two elevations in a free space calculation at a frequency in the millimeter wave region. Mesh generation times and analysis run times will be compared to determine total production time tradeoffs. The key questions are: will the use of additional primitives be faster and if so, are the results comparable? If this investigation yields significant results the next question would be, can these results be scaled to full size vehicle targets on the 100,000 to 1,000,000 element range?

INTRODUCTION

Creation of a suitable closed surface model from CAD models employing combinational solid geometry and parasolids is a time consuming task. Some of these models have a different primary purpose from that of observables. For example, determining space reservation and interference, as well as weights of modeled objects may help other design considerations, but they are excess baggage in observables analysis. Geometric entities do not always have concise intersections and intersection edges have to be derived, sometimes with great difficulty.

Observables Analysis is concerned with the exterior surface only. For example, (RTS) requires closed surface geometry defined by points in 3 dimensions and elements defined by these points. The process was initially refined about five years ago with the development of Pre-RTS, a CAD (really Finite Element) to POLY file (RTS Input file) geometry translator. It takes a PATRAN Neutral File and produces a POLY.DAT file. Naval Sea Systems Command uses Intergraph's Finite Element Modeler (I/FEM) to produce a PATRAN Neutral File. Standard practice was to create mesh, because much source geometry was just a wire frame definition and the desired output was a finite element mesh. Because there was no need to create a CAD Model only the geometry needed to create a mesh was constructed. However, this process is very time consuming so faster, more accurate approaches are always being investigated and implemented into the modeling process. Final validity in the X-Windows Radar Analysis Tool (XRAT) and the successful completion of XPT and TDI (two preparation modules of RTS) are the driving factors.

When RHINOCEROS (RHINO), a Non-Uniform Rational B-Spline (NURBS) modeling environment, with a user-friendly fast automesher, which uses tri and quad elements, was proposed as a way to automate the process because it had a more versitile IGES geometry interpreter, it was decided that its use in the process should be investigated. The following sections are a description of the modeling and analysis process with a final section of observations on the pros and cons of each approach.

MODELING
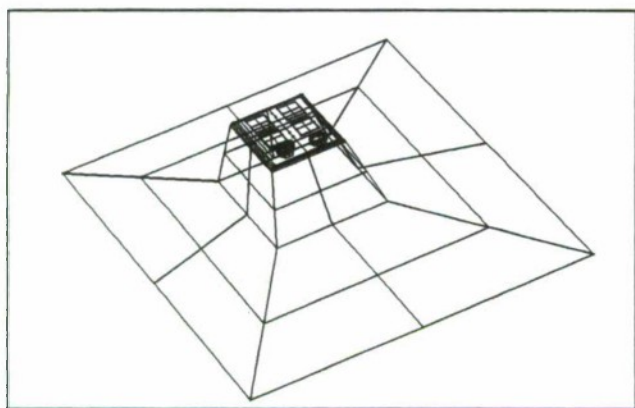
Figure 1

The basic model, shown in Figure 1, started as a rectangular cube. The sides were trimmed to be 20 degrees positively sloped, or so their surface normals were 20 degrees above the horizon. The rear face was sloped positively 10 degrees and the front face was sloped 30 degrees. The cube was then trimmed again to round off three of the top edges (sides and back).

A mismatch resulted at the rear corners because the sides and back had different slopes. Consequently, the trimming of the back top edge had a longer arc length than the sides, introducing a splinter triangle element to make up the difference (the gray highlighted element as shown in Figure 2.) and leading to meshing errors with the auto mesh.
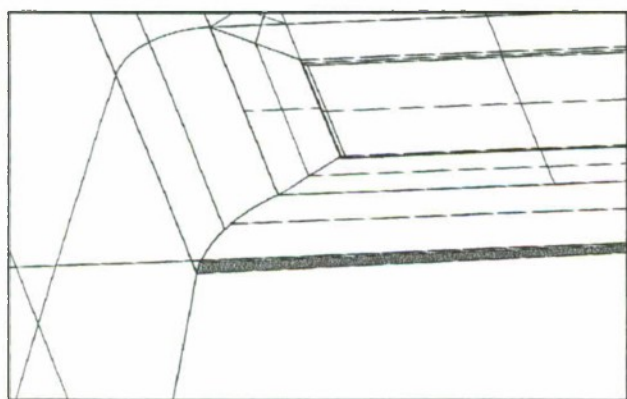
Figure 2.

In the interests of time a compromise was done. The side arcs redefined the back surfaces and the side flat faces include the splinter in a new larger quad. This example shows that, even with the latest methods of automatically extracting surfaces and automeshing, the underlying CAD geometry sometimes needs to be fixed. This raises the question of how much one can trust these automatic methods.

Four different types of hatch entities were added to the model. In clockwise order starting at the lower left in Figure 3: a stand-off door with a quarter wave length stand-off crack, a raised door with a double curved conforming edge, a depressed door with a single curved edge, and a conformal door with a quarter wave length wide and deep gap. Next the front face was then extended to include two polygonal entities, the left one with both positively and negatively sloped sides and the right one with just negatively sloped sides. The new front face did not have a rounded top edge. Fastener features were then added in three sizes; quarter, half and full wavelength and two major types; hemi-spherical, and conical. Finally there were two weld bead entities, one with a half circle cross section and the other a flatter arc where the length of the mid-radius is half that of the end points.
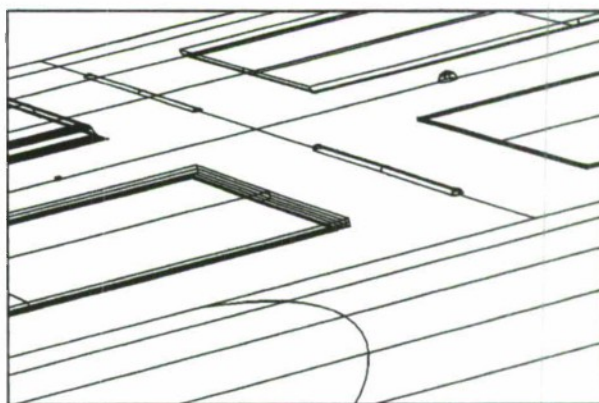
Figure 3

Two view port shapes were then added to the front of the model as shown in Figure 4. Additional horizontal surfaces were added to the bottom edge to replicate interaction of the model with the ground on which it sits.
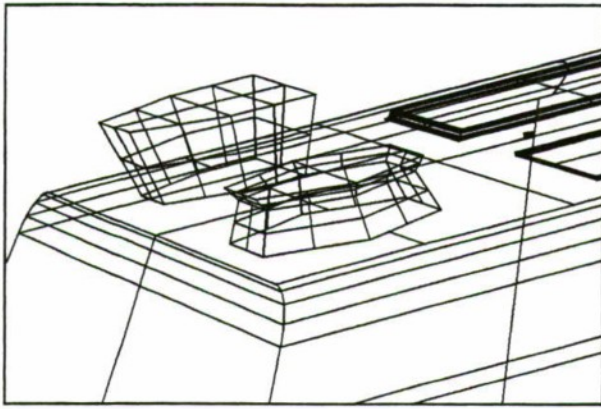
Figure 4

A traditional Finite Element Model was then produced in I/FEM using additional primitives, partial cylinders and cones. A PATRAN neutral file was output from I/FEM. Pre-RTS checked and translated into a POLY.DAT file. Finally, it was checked in RTS by XPT and TDI modules, multibounce sequences were generated by Multi-Scatterer Generator (MSG) and analyzed by SCP modules in full azimuth cuts at 32, and 60 degree elevations.

While the first model was undergoing analysis, the second model was being produced. An IGES model was first exported to RHINO. All the surface entities were then joined together to make a complete closed surface. This was meshed and exported as a Virtual Reality Modeling Language (VRML) model. This VRML model was translated into a PATRAN Neutral File and brought back into I/FEM to check for validity. Of note, the automeshed model was found to have a few problems. First, the initial mesh density was too high causing the creation of numerous splinter shaped tri elements. Second, even though RHINO seemed to handle the surface geometry as one complete surface when it was meshed, there were duplicate nodes at surface entity edges. After several attempts at lowering the mesh density, only one splinter tri remained which was remeshed manually.

Taking all things into account, the use of RHINO to create the mesh was faster and easier. There is somewhat of a learning curve, but far less than with most CAD or 3D drawing packages. It is by no means a FEM environment, but it is so fast that doing things over doesn't take long. In the final analysis, automeshing has gotten smart enough to be faster and more efficient than manual meshing methods.

ANALYSIS

The analysis phase is where the advantages of additional elements supported by RTS begins to pay off. RTS is made up of several modules divided into three groups: target file preparation, signature calculation and post processing. Even with a valid geometry definition, there is still a great deal of runtime before any calculations can be done. The multi-bounce scatterers must be determined using the MSG module. On full-scale models this can take weeks of runtime. The advent of fast workstations has reduced this somewhat, but RTS is always trying to reduce runtime. It is a central feature of its design. Comparison between the two models demonstrate this in Table 1:

Model Type	Manual Mesh	Auto Mesh
Element Count	437	5689
Max Freq.	100 Ghz	100 Ghz
Threshold	-80dBsm	-30dBsm
# of Bounces	17	8
MB Scatterers	41812	65422
MSG Runtime	14.5 hours	1.1 hours

Table I

There seems to be a vast difference in runtimes but the automeshed model ran for several hours with a threshold of −50 dBsm and did not even get through three bounces, an indication that lowering the threshold would make the MSG runtime too long. So it was raised back to −30 dBsm for this analysis. The manually meshed model ran through several MSG calculations in a matter of minutes. This final run was left to run overnight for the evaluation of 17 bounces. Note that the multi-bounce sequence counts are about the same magnitude.

Results are comparable in that the peaks overlay quite well and the mean and median values are well within acceptable limits of each other. There are additional peaks in the automeshed model probably caused by the facetization of the curved top edges. The manual meshed model seemed to have lower troughs between the peaks.

The Signature Calculation Process Runtime Table follows:

Model Type	Manual Mesh	Auto Mesh
1800 calc points	278.35 sec	1294.29 sec
450 calc points	80.2 sec	330.21 sec

Table 2

These runtimes indicate the advantage of using additional primitives in RTS. Both types of models were completed in a matter of minutes, and have been demonstrated to be scalable using RTS primitives. A 90k

viewable element model with an 8 bounce sequence file completed a single elevation 360 degree azimuth in less than 18 hours or overnight on a single processor DEC Alpha PC. If faster times are required, multiple machines could be used and the runtimes reduced accordingly. Additional processors could be used to make the automesh models useful too, but the process would be slower than analysis of a manual meshed model.

SUMMARY

The results indicate that although automeshing is a faster method to create a valid closed surface mesh, the increases in analysis runtime do not balance out the saved time. Also, there are trade-offs: more care must be taken to ensure that the surface to be automeshed is correct, that no cracks or discontinuities exist, that it is all joined together properly, and duplicate nodes removed. Automeshing may be practical for a small object or group of objects, but scalability for larger, more complex models does not seem to be practical. The resulting increase in element count automeshing incurs increases run times for MSG and SCP modules of RTS, making the use of automeshing to support design impractical.

Manually meshed models still enjoy the advantage of keeping element counts low. With the development of addtional primitives like 6 and 8 noded doubly curved elements, RTS should be able to handle rapid analysis in support of the design process. RTS is continuously enhancing its capabilities, particularly in the area of analysis. The upcoming ability to display results as real color levels on displayed models in XRAT is due in the next release. Work is ongoing to include cylinders and cones in multibounce sequences.

ACKNOWLEDGEMENTS

Multi-Spectral Characterization of Complex Ground Targets

Clifford D. Geohagan
James L. Cook
Samuel D. McKenzie
Scott R. Vechinski

Science Applications International Corporation (SAIC)
6725 Odyssey Drive
Huntsville, AL 35806

ABSTRACT

In the production of CAD models suitable for computational electromagnetics analysis, it is important to determine the required level-of-detail (LOD) and level-of-resolution (LOR) of the models for accurate signature synthesis. To help determine these requirements for high-frequency analysis, models of five ground targets were generated with varying levels of detail and resolution. Synthetic images were produced at numerous elevation and azimuth angles from scattering centers and were then compared to measured SAR images using an automated, quantitative image comparison technique that includes image segmentation, resampling, scaling, and filtering. Error metrics are presented for a variety of cases, and a simple ATR application is also shown for the varying LODs and LORs

At lower frequencies, the Fast Illinois Solver Code (FISC), an industrial-strength code that employs the method of moments and the Multilevel Fast-Multipole Algorithm (MLFMA), was used to solve the scattering problem for complex targets with up to hundreds of thousands of unknowns. The characterization of complex ground targets for various LOD's using FISC at frequencies of 1 GHz and below will be presented. Insight into FISC and the use of the MLFMA for geometries with larger numbers of unknowns will be provided.

INTRODUCTION

Science Applications International Corporation (SAIC) supports the Virtual Target Program with the development and characterization of physically realistic computer-aided design (CAD) models, or virtual targets. The program is sponsored by the U.S. Army Simulation, Training and Instrumentation Command (STRICOM), Program Manager for Instrumentation, Targets & Threat Simulators (PM ITTS), Targets Management Office (TMO). TMO is addressing the requirements for simulation-based acquisition, test and evaluation, and training by producing virtual targets that possess enough fidelity to synthetically replicate the electromagnetic characteristics of real threat targets across a broad frequency spectrum. However, the question of "how much fidelity is enough?" has never been addressed. The required fidelity is dependent on a number factors to include the frequency of interest, required data resolution, and the simulation software used to generate the synthetic data. The study described here provided an initial assessment of the required virtual target fidelity for specific, but common applications that span the radio frequency (RF) spectrum from 150 MHz to 10 GHz [1].

BACKGROUND

An important aspect of a weapon system assessment is its radar cross section (RCS). It has been demonstrated that accurate, cost effective RCS analysis can be performed on threat systems using high-fidelity CAD models and computational electromagnetic (CEM) codes to create synthetic radar signatures. To date, most of the emphasis has been placed on high-frequency signature analysis including frequencies in the S-band region or higher. However, an area of growing interest is the low-frequency (LF) response of these systems, which can vary significantly from the high-frequency response.

Most CEM codes can be characterized as using either a "low-frequency" or "high-frequency" approach. In this context, low and high frequency refer to the electrical size (*i.e.*, the size in wavelengths) of the object of interest. The CEM algorithms in these two regimes are quite different.

For example, asymptotic techniques work well in the "high-frequency" regime where targets are electrically large. For these size targets the scattering mechanism is primarily a localized phenomenon where the induced currents at a given location are not significantly dependent on induced currents elsewhere on the target.

The resulting CEM solution, with methods such as the Physical Theory of Diffraction (PTD) and the Geometrical Theory of Diffraction (GTD), is thus decoupled. The required run time is more affected by the number of required aspect samples rather than by the number of frequency samples.

Conversely, codes employing Method of Moments (MoM), Finite Element, Finite Volume, and other such algorithms provide "exact" solutions to the scattering problem but are typically limited by memory requirements and computation time. Therefore, these type codes are generally used for electrically smaller targets. Scattering from these size targets is primarily a global phenomenon where the induced currents at a given location can significantly depend on the induced currents elsewhere on the target. MoM techniques require the solution of coupled integro-differential equations where the required run time, in contrast to the high-frequency methods, is more affected by the number of required frequency samples than by the number of aspect samples.

Facetization and level-of-resolution create even more distinction between "low-frequency" techniques and "high-frequency" techniques. For high-frequency analysis, facetization requirements exist in order to accurately represent curvature with respect to wavelength. For low-frequency analysis, facetization requirements are needed to accurately represent the variation of the induced currents and usually have strict limits on how much facetization is needed based on the operating wavelength and the target size.

All of these parameters have a direct impact on the required fidelity of the virtual target. At higher frequencies, small details become more important as the frequency increases, while at lower frequencies, the overall dimension and shape of the major features have greater influence in obtaining accurate RCS characterizations.

HIGH FREQUENCY ANALYSIS

In the generation of high-fidelity CAD models, two key issues arise. First, there must be a determination by the modelers concerning whether or not an individual feature should be modeled. This criterion is referred to as feature existence or level of detail (LOD). Second, when a feature is modeled, the modelers must determine what facet density to use to properly represent the feature curvature. This criterion is the level of resolution (LOR) parameter.

The answer to these issues is based upon the ultimate use of the models and of any data generated from the models. In the context of this study, the models are to be used for synthetic electromagnetic signature generation at X-band. The conclusions reached in this analysis are therefore dependent on the application for which the models and any subsequently generated data are used.

The geometric modeling criteria used during the generation of a model are dependent on a number of conditions. Since electromagnetic scattering inherently depends on the size of features in terms of wavelengths, the frequency of operation of the radar is an important factor to consider. The form of the data, whether it be total RCS, range profiles, or synthetic aperture radar (SAR) imagery also plays a part. Resolution parameters for these data formats also play a significant role. Finally, the processing steps used to operate on the data also affect the conclusions.

This study was designed to be as comprehensive as possible with the current state-of-the-art in CAD modeling and electromagnetic signature generation. Five ground targets, the BTR-70, M-3, M-35, T-72, and ZSU-23/4, were chosen for this study. As an example, a photograph of the ZSU-23/4 used during the measurement process and an image of the baseline virtual model are shown in **Figure 1**. The highly accurate, physically representative geometric models of these targets were configured according to different LODs and LORs. Signatures were at numerous elevation and azimuth
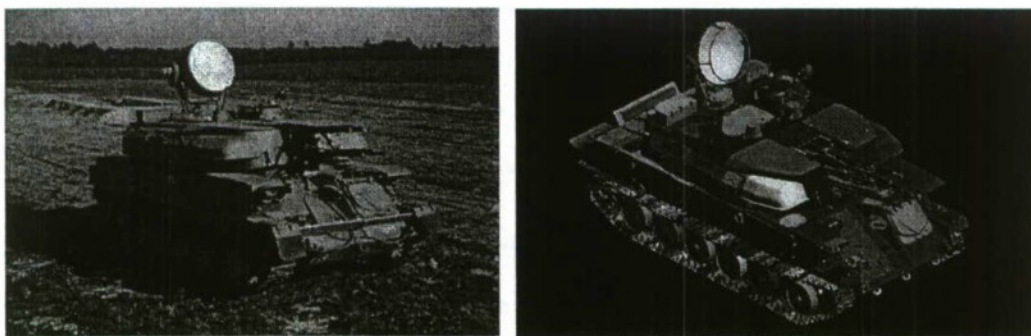

Figure 1. ZSU-23/4 and corresponding CAD model

angles. All of this data was compared to measured SAR imagery using various signal processing techniques.

Data Generation Tools

As noted above, most CEM codes can be characterized as using either a "low-frequency" or "high-frequency" approach. The parametric analysis presented in this report examines the fidelity requirements using both approaches. For the high-frequency analysis, all of the synthetic signature generation was performed using the Xpatch (version 2.7) suite of prediction software tools.

Xpatch is used predominantly in high frequency RCS computations where the RCS can be calculated using asymptotic approximation techniques such as Physical Optics (PO) and PTD. In conjunction with PO/PTD, Xpatch determines where energy is incident on the target and how energy from one point on the target interacts with other points (multi-bounce) by employing a ray tracing technique called "Shooting and Bouncing Rays" (SBR). In addition to multi-bounce, SBR also provides Xpatch with a convenient way of determining the shadowing of one portion of the target by another.

Xpatch does not, however, account for higher order scattering mechanisms such as traveling or creeping waves. These higher order effects provide little contribution to the RCS at the higher frequencies, but become major contributors as the targets electrical size decreases. In addition, the modeling of non-conducting materials poses difficulties for all high-frequency asymptotic computational electromagnetic analysis codes, including Xpatch. Fortunately, the RCS of most ground targets is dominated by the specular returns of large metallic features, such as hulls and turrets. It has been demonstrated through measured-to-synthetic data comparisons that removal of non-metallic materials does not have an appreciable effect on the predicted RCS data for these targets. Therefore, to minimize run times all non-metallic materials were removed from the targets included in this task.

Level-of-Detail Configurations

Electromagnetic scattering is significantly dependent on the frequency of the radar. Targets, and the individual features composing the target, produce different scattering mechanisms depending on the size of the features with respect to the associated wavelength. Furthermore, the path-length relationships between scattering locations, and thus the complex summation of the scattered fields, is affected by the frequency. The task at hand is to determine what LOD is necessary, at the particular frequency of operation, for accurate signature generation.

Another significant factor in determining the required LOD for a given accuracy is the format and resolution of the measured data that the synthetic data is required to replicate. For this analysis, the measured data is in the form of SAR data with approximately 8 by 8 inch pixels.

Some additional qualifications must be made, however. Previous studies have indicated that although the scattering values from electrically-small features is somewhat suspect, these contributions can be vital to the target signature if these scattering mechanisms are the primary contributors in a given range bin (for range profiles) or in a given range/cross-range location (for SAR images) [2]. This phenomenon is more prevalent for air-targets, such as missiles, than for ground targets.

A number of different criteria, such as surface area, volume, and maximum dimension, may be used to generate the LOD variations. Surface area was ultimately chosen because it is generally more indicative of the level of scattering than the other criteria.

Based upon the results of a previous study that identified the signature differences for ground targets with respect to synthetic baseline data [3], LOD levels for the current study were set at surface areas of 5, 10, 50, 100, 250, 500, and 1000 square inches. Thus, the 5 square inch LOD configuration, designated by LOD0005, was generated by removing all of the features with a surface area less than 5 square inches from the baseline model. This process is quickly performed in SAIC's *ModelMan* software environment, which has a feature that automatically identifies which parts should be toggled off based upon a user-defined setting of the computed surface area. As an example, the LOD0100 variation is shown in **Figure 2**, and the parts that were removed from the baseline model in order to generate this variation are shown in **Figure 3**.

Level-of-Resolution Configurations

For the accurate generation of electromagnetic scattering at high-frequencies, it is also necessary to properly represent the surface curvature in the model. The term level-of-resolution is used to identify the different amounts of facetization on curved surfaces.

The accuracy in representing a curved surface is dependent on the radius of curvature. The smaller the radius of curvature, the greater the necessary facetization to accurately represent the curved surface. At one extreme is a flat plate, which has an infinite radius of curvature. For high-frequency scattering models, only two-triangular facets are often used to model flat plates.
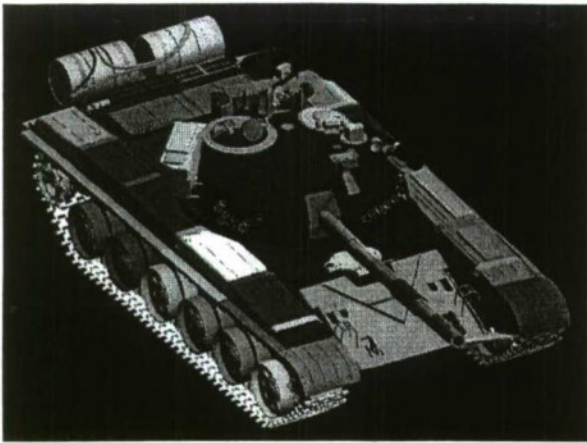
Figure 2. LOD0100 variation of the T-72

Figure 3. Parts removed for LOD0100

However, the local radius of curvature (which in general is a function of position on the surface) is a property that is often not readily available in the model. Some facetization algorithms can use a parameter related to radius of curvature, such as chord deviation tolerance, however this is neither universal nor is it the only parameter within meshing algorithms. Thus, an appropriate criterion for LOR variation is more elusive than for LOD variation.

After considerable effort in using various combinations of available parameters to define an LOR criterion, it was noted that the original facetizations of the target features inherently contained a reasonable consistency with respect to facetization. This consistency occurred because the CAD modelers instinctively adjusted the meshing parameters to produce features with similar facetization levels with respect to curvature as noted by a visual inspection. Thus, it was determined to used percentage reductions in the facet count of the baseline feature as the LOR criterion. The LOR reduction levels were set at

87.5%, 75.0%, 62.5%, 50.0%, and 37.5%. The model variations are designated by the names LOR875, LOR750, LOR625, LOR500, and LOR375, respectively.

Due to previous experience in reducing the facetization level [4], where relatively large errors in the signatures were noted, a very strict visual inspection of the new surfaces was performed. If the resulting surface was noticeably displaced at a given reduction level, the facet reduction for the part was stopped at the previous step.

Thus, the five LOR configurations were generated for each of the five targets using this procedure. As an example, the LOR variations for the T-72 turret are shown in **Figures 4 and 5**.
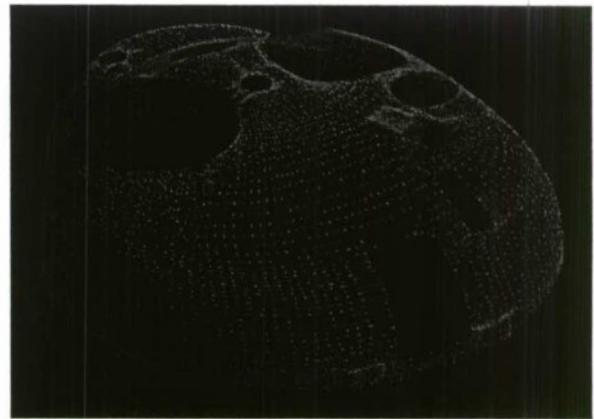
Figure 4. Baseline for the T-72 turret

Figure 5. The LOR375 variation

Analysis Methodology

The methodology used to generate the SAR validation and comparative metrics that were implemented in this

59

study is now discussed. This process has the desired properties of being automated, repeatable, and modular.

Each measured image is operated on in turn. In order to replicate a typical image processing system, the image is first converted (or mapped) to an integer-valued one with pixel values ranging from 0 to 255. The measured image is then segmented to remove the target from the background noise. As the measured data for this study did not have uniform pixel sizes from image-to-image, the segmented measured image is then resampled to provide an image with the same pixel dimensions as the synthetic data. Finally, this result is then low-pass filtered.

Because the XpatchT data is in complex scattering matrix form, the complex-valued candidate synthetic SAR image is first converted to a real-valued image with linear units. This image is then converted to dBsm via a logarithmic transformation, and then to an integer valued one. The synthetic image is then low-pass filtered. In order to account for offsets between the measured and synthetic data sets, the measured and synthetic images are then equalized by equating the means of the pixel magnitudes.

Before applying the metrics, however, the two images are shifted over various combinations in both the horizontal and vertical directions. This process accounts for the lack of registration between the measured and modeled images. Thus, the metric is used to evaluate the comparison at the various shifts, and the best comparison over all possible shifts is retained as the best error metric for that candidate image.

The remaining candidate synthetic images are then each evaluated with the measured image. The candidate synthetic image with the smallest error metric, *i.e.*, the best matching score, is identified as the best match. The process is then repeated for each measured image. After looping over all the measured images in an elevation set, mean metrics are computed by averaging the metrics over azimuth.

High Frequency Analysis Results

As previously discussed, the LOD variation study consisted of a baseline geometry model for each of the five targets (BTR-70, M-3, M-35, T-72, and ZSU-23/4), plus seven LOD variations for each target. Scattering centers were produced on each of these 40 models at a variety of elevation and azimuth angles. SAR images were then computed at 1° steps in azimuth at the three elevation angles of interest (17°, 29°, and 46°). The images in each 5° azimuth window were averaged in order to produce an image template at 5° steps in azimuth.

These synthetic image templates were then compared to the measured data.

The correlation error for each measured-to-synthetic comparison was computed, and the lowest error was stored for each measured image. The average CE over all measured azimuths was then determined at each elevation angle. These average correlation errors are thus the average errors for the approximately 275 measured images at each elevation angle for each target.

The original criterion for producing the LOD variations was surface area which has units of length squared. By inspection of the data, it was determined that it was more useful to use a linear length criterion for plotting purposes. Thus, the concept of equivalent side length (ESL) was generated. The ESL is defined to be the length of a side of a cube with the same surface area as the modeled part, or

$$ESL = \sqrt{\frac{SA}{6}},$$

where *SA* is surface area of the equivalent cube. The ESL axis is also scaled in terms of wavelength. This formulation allows the analyst to think in terms of the removal of parts based upon the length in wavelengths, instead of the surface area of the part, which is a less intuitive parameter.

The individual CEs at each azimuth angle (for each LOD variation at each elevation angle) were also compared to each other to determine which modeled target produced the lowest CE. These predictions over the approximately 275 measured images at each elevation angle (for each target) were then used to create a confusion matrix which identifies the percentage of cases in which the algorithm predicted that each target was present. **Figure 6** illustrates the overall average of correct identification for all models over each elevation.

As seen in the figure, the percentage of correct identification begins to degrade significantly as the size of the parts removed becomes greater that 3 ESL, or 3 wavelengths on a side. This general trend was noted in all models.

The LOR variation study did not produce similar trends in the metric scores. In retrospect, the use of only a small number of scattering features and the strict refacetization criterion resulted in no noticeable change in the CEs and the corresponding identification percentages. These LOR reductions reduced the overall facet count of these models by approximately 20% for the lowest LOR level.
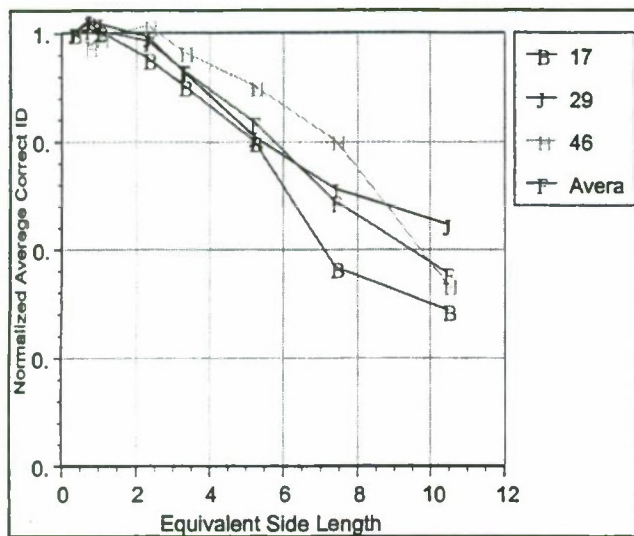
Figure 6. LOD study results. Average correct identification percentages for all targets versus the equivalent side length of the components removed.

LOW FREQUENCY ANALYSIS

For the low-frequency analysis, the synthetic signature generation was performed using the Fast Illinois Solver Code (FISC). FISC employs the method of moments (MoM) technique and the Multi-Level Fast Multipole Algorithm (MLFMA) to solve for the induced currents on the surface of the target which is described by a triangular facetted mesh. The fast multipole method is an efficient way to perform matrix-vector mulltplies whereby the field at each of the vertices of the triangular mesh, due to every other point on the target, is calculated for all the vertices comprising the target surface definition. Unlike high frequencies, low-frequency electromagnetic energy penetrates thin non-metallic materials making them appear to be totally transparent. Therefore, for consistency and to keep the run times at a minimum, all non-metallic parts were also removed from the low-frequency targets prior to generating the synthetic RCS data.

The low-frequency parametric fidelity analysis was accomplished by first creating FISC compliant CAD models from two baseline high-fidelity virtual targets. Three different baseline FISC models were created to support the mesh requirements across the low-frequency spectrum. The fidelity of these baseline targets was then varied and synthetic RCS data was generated for each fidelity variation. Unfortunately, no useable measured data was found to perform measured-to-synthetic data comparisons. SAIC did make several requests for measured data and actually received low frequency measured data collected in support of the Defense

Advanced Research Projects Agencies Foliage Penetration (FOPEN) program. However, the data had been preprocessed and was unusable for this effort. Therefore, each fidelity variant was compared to its baseline FISC compliant model for the specific frequency of interest.

CAD Model Design

The facet mesh requirement presents a difficult challenge. FISC requires the model to have matched edges due to how the MoM technique determines the electromagnetic interactions. To accurately represent a complex geometric model, the curvature of the surfaces of the model must be maintained. Without the representation of curvature, the model visually appears blocky and can produce anomalous RCS values. Since the wavelength in the low-frequency region is large, smaller parts on the model can be eliminated, reducing the complexity of the model. The model must be fitted with a uniform mesh that is specific and "optimum" for the frequency being tested. In other words, for each frequency band within the "low-frequency" region, a new model must be developed with a mesh that is specific to that frequency band. If the model being used by FISC is determined to have a facet maximum edge length greater than the user specified accuracy requirement, then FISC will automatically re-discritize the surface of the model. It does this by sub-dividing the original facets to the point that the maximum edge length parameter is satisfied. Allowing FISC to perform this action may only be valid if the model is used within a small frequency deviation from what the model was built because the curvature of the surfaces is not updated as it remeshes the model.

Building a new complex model for each frequency band is a non-trivial problem. If the model was originally built using a single, advanced, solid-modeling computer-aided design (CAD) package, e.g., AutoCAD, the surfaces can be welded together and the resultant solid model can be exported to another program to generate a uniform mesh. If the model exists only in a facet representation, then the problem becomes complicated in that the surfaces cannot be extracted directly from the model.

To approximate the surface of the model in an automated fashion, a concept was developed that implements volumetric techniques. To accomplish this, the space containing the input geometry is voxelized, or quantized, into cubes whose size is determined by the input edge length. This voxelization is simply a structured point set bounding the input geometry. The vertices of the voxelized space are assigned values that record the distance from that voxel vertex to the underlying input geometry. All distances are positive. Because of the quantization of space, you cannot record signed distances.

61

This leads to separation in parts having thin or small areas. Using the "Marching Cubes" algorithm, an isosurface is extracted from the structured point set. A contour value must be specified. Since all scalars are positive, you cannot extract the zero surface, which corresponds closest to the underlying geometry. Instead, you must specify a contour value of at least the size of a single voxel. This generates geometry that is offset from the underlying geometry. A positive side effect of this is that disjoint components in close proximity (no more than the offset apart) to one another in the input geometry are combined. The "Marching Cubes" algorithm tends to create some facets with short edges or with an unacceptable aspect ratio. To overcome this, the geometry's point coordinates are adjusted using Laplacian smoothing. This relaxes the mesh and makes the facets better shaped. Using the ModelMan software, you can check the compliancy of the resulting FISC model. ModelMan will report min edge length, max edge length, max/min edge length ratio, min aspect ratio, max aspect ratio, and average aspect ratio. To use the program, the user needs only to have a facet representation of the geometry. The user may pick the maximum edge length as input to the code. This code, called FISCulator, was used to develop the 1000 MHz models of the ZSU-23/4 and the T-72 [5].

FISC Parameters

As with any code, production parameters can be set to balance the quality of the results versus computer resource requirements. In this case, relaxing the accuracy of the EM solution will lower resource requirements, which will allow for a realizable solution. Two of the FISC parameters of most importance are the maximum edge length and the solution method. The maximum edge length is defined such that a particular fraction of a wavelength will be present across the edges of the facets, i.e., many edges comprise one wavelength. The maximum edge length must be chosen correctly to obtain accurate RCS values. The consequence of this parameter is twofold; first, the geometric model must be conformed to this specification, which becomes prohibitive at higher frequencies; second, the CPU and memory requirements are increased because of the number of unknowns developed due to a higher number of facets on the model. If the number of wavelengths across the edge is too large, then the accuracy of the RCS will degrade. As the number of wavelengths becomes small, the memory requirements rise, and can realistically approach 3 GB or more. For the purposes of this study, it was determined that a maximum edge length value of 0.2 wavelengths would be sufficient to maintain a reasonable RCS while controlling memory requirements [6].

The solution method involves the method by which the matrix developed by the code is solved. Solution techniques include an LU Decomposition method, an iterative method using the full matrix, and the Multilevel Fast-Multipole Algorithm (MLFMA) [7]. Using LUD, the full matrix is used and the entire data matrix must be stored in memory. For smaller targets, this method is preferred because the matrix is solved exactly. Yet, LUD is prohibitive for large, complex targets because of the memory requirements introduced by higher numbers of unknowns. The iterative method, which is also useful for small targets, performs an iterative solution to the matrix, providing an approximation to the solution. This method saves run time and memory requirements over LUD, but is still only realistically used for small targets.

The MLFMA is an iterative solution of the MoM matrix utilizing subsections of the full matrix to solve the problem. "The fast multipole method is an efficient way to perform matrix-vector multiplies whereby the field at each particle due to every other particle is calculated for all particles in an ensemble of N particles. Ordinarily, this would require $O(N^2)$ calculations. With the fast multipole method, this can be reduced to $O(N)$ or $O(NlogN)$ depending on the spatial distribution density of the particles and implementation" [7]. In other words, the MLFMA helps to reduce the memory requirements of the code by reducing the complexity of the matrix at each step. The tradeoff is a controllable but induced error in the approximation to the calculated RCS.

Model Generation and Variation

As described earlier, the models needed for this task were generated using several different techniques. To generate a model for 1000 MHz, the FISCulator was used to assimilate the available high-fidelity high-frequency model for FISC geometry compliance. The resulting models of the ZSU-23/4 and the T-72 are illustrated in **Figure 7** and **Figure 8**, respectively.

Figure 7. ZSU-23/4 1000 MHz Solid Model 273,354 facets ~410,000 Unknowns

Figure 8. T-72 1000 MHz Solid Model
402,624 facets   ~604,000 Unknowns

These two models were used to produce signatures at 1000 MHz using the Space and Missile Defense Command (SMDC) Advanced Research Center (ARC) High Performance Computer (HPC). Because of past computer resource limitations, the production of signatures using FISC at frequencies at and above 1000 MHz has been mostly limited to smooth geometries and bodies-of-revolution. With the development of FISCulator, the more complex models of the ZSU and the T-72 were realized.

To generate models for 150 MHz and 450 MHz for the ZSU, the original CAD geometry file was used. Since the ZSU was originally modeled using a single CAD package, this process was straightforward. Small parts were removed and a solid model was created within the CAD environment. The resulting geometry was read into a third party meshing tool called MSC/NASTRAN, which created the FISC compliant mech.

In addition to the baseline model that contained all the parts deemed appropriate for the frequency band, three level-of-detail (LOD) models were developed. The signatures of these models were then compared to the baseline model. The baseline and LOD variations for the ZSU are shown in **Figure 9**. LOD 1 is configured as the baseline minus the small parts on the turret. LOD 2 is configured as LOD 1 minus the guns. LOD 3 is LOD 2 minus the radar assembly.

Development of the 150 MHz and 450 MHz models of the T-72 was more complicated. Since this target was not originally built using a single CAD package, parts of the model had to be rebuilt using a common CAD environment. The MSC/NASTRAN was used to create the mesh. Although time consuming, the baseline model and three LOD variations were and are illustrated in
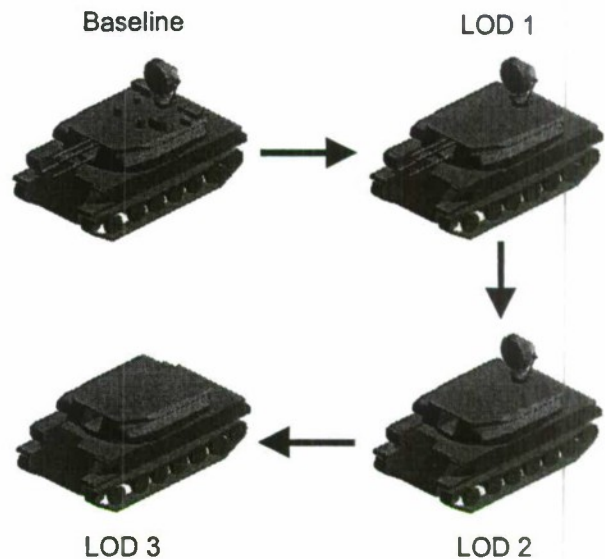
Figure 9. LOD Variations for the ZSU-23/4

Figure 10. The LOD variations of the T-72 are similar to that of the ZSU. LOD 1 is configured as the baseline minus the small parts on the turret, including the smoke canisters and the toolboxes. LOD 2 is configured as LOD 1 minus the main gun. LOD 3 is configured as LOD 2 minus the fuel barrels.
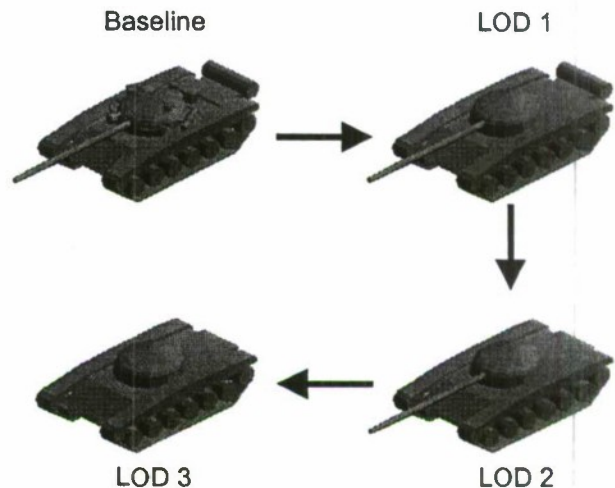
Figure 10. LOD Variations for the T-72

LOW FREQUENCY RESULTS

(Plots of the total RCS of the ZSU and the T-72 are not shown due to the need for grayscale images in the documentation. Please refer to Reference 1 for these

63

plots.) As stated earlier, the scattering mechanisms at low-frequencies are global phenomenon in that every part on the target interacts with every other part. The scattering, especially at resonance where the target length approaches one wavelength, is an integration of the energies from all areas of the target that are either illuminated by the impending electromagnetic wave, or are influenced by induced traveling waves on the surface of the target. Discontinuities in the current flow will promote higher levels of scattering. Overall, it is found that large structural components of the target tend to impact the overall radar cross section more that those parts that are either smaller in size or are only large in one dimension. The definition of large and small in this context refers to the parts electrical size in wavelengths. Thus, as expected, removing the large radar dish from the turret of the ZSU has a dramatic impact on the RCS at nose-on and rear-on. In this case, the bulkiness and size of the radar dish provides a significant contribution to the scattering, especially at those angles. For the T-72, at 150 MHz the fuel barrels on the rear of the T-72 appear to have a notable effect on the RCS at rear-on aspects where as the small parts on the turret do not. Once again, since the fuel barrels are rather large objects when compared to the parts on the turret, it is expected that the scattering from the barrels will have greater affect on the overall RCS. Since the wavelength at 450 MHz is 1/3 of the wavelength at 150 MHz, the definition of small parts versus large parts changes. At this higher frequency, the parts that are on the turret of the T-72 that were taken off for the LODs have an impact on the overall RCS.

Thus, variations in the total RCS of a target at low-frequencies is dependant upon the location and structural size of its individual parts and the definition of those parts as large or small based on the frequency. Parts that have an electrically large surface area, e.g., the radar dish on the ZSU, will have a significant impact on the total RCS. Those parts that are electrically small or only have a greater size in one dimension will not have a significant impact on the total RCS.

At 1000 MHz, the components on the model become more important to the overall scattering from the target. Since the wavelength is reduced significantly, the region of operation is in between that of low-frequency and high-frequency. To examine the low-frequency approach to solving the scattering problem, FISC was used to develop signatures from the ZSU and T-72 at 1000 MHz. **Figure 11** illustrates the current distribution on the ZSU at nose-on.
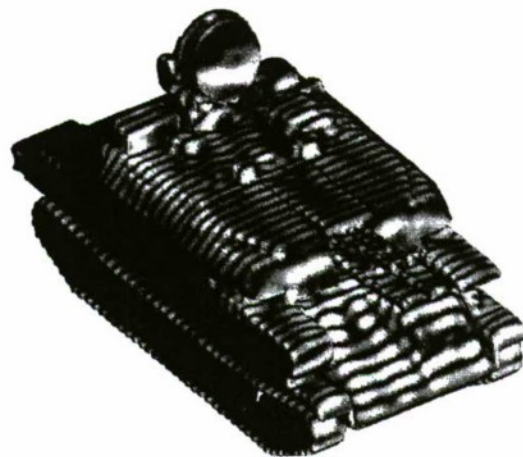
Figure 11. Current Distribution on the ZSU-23/4 at 1000 MHz, illuminated at 10° elevation and 0° azimuth

Figure 12 illustrates the current distribution on the T-72 at 45° azimuth. The peaks and valleys of the current distribution can be seen, providing a feel for the currents that are induced on the target at 1000 MHz. Note that the wheels on the T-72 are illuminated at 45° azimuth while those of the ZSU at 0° azimuth are not.
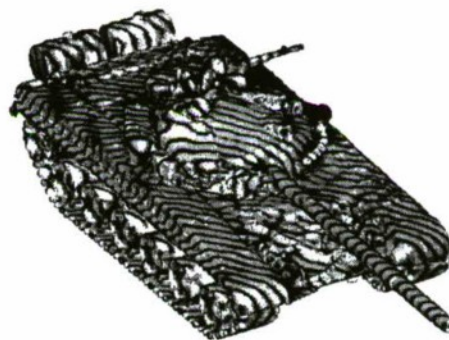
Figure 12. Current Distribution on the T-72 at 1000 MHz, illuminated at 10° elevation and 45° azimuth

FISC Runtimes and Memory Requirements

The majority of all low-frequency production for this task was accomplished using the ARC's HPC. **Table 1** provides a breakdown of the runtimes that were incurred for each system and its LODs. The runtime numbers are in hours. Of note is the difference between the number of hours needed for each frequency level. Since a greater number of facets are needed to accurately represent the target as the frequency increases, the number of unknowns also increases.

Run Times (hours)	Baseline	LOD 1	LOD 2	LOD 3
ZSU 150 MHz	79.50	81.73	68.62	64.04
ZSU 450 MHz	460.70	350.37	343.59	293.49
ZSU 1000 MHz	546.40			
T-72 150 MHz	61.44	37.29	28.60	21.69
T-72 450 MHz	291.52	420.74	314.88	447.40
T-72 1000 MHz	737.49			

Table 1. Run Times for the ZSU-23/4
and T-72

Table 2 illustrates the memory requirements for each system. The bold numbers indicate the memory requirement using FISC v1.0. The rest of the values indicate the memory requirement when using the updated FISC v1.3. As before with run times, as the operating frequency is increased, a greater number of facets are needed, which in turn increases the amount of required memory needed to complete the simulation.

Memory Requirement (MB)	Baseline	LOD 1	LOD 2	LOD 3
ZSU 150 MHz	**73.7**	68.6	61.2	54.7
ZSU 450 MHz	**325.9**	251.5	234.9	222.6
ZSU 1000 MHz	2245.1			
T-72 150 MHz	109.8	73.9	55.2	27.0
T-72 450 MHz	126.9	175.5	135.7	251.7
T-72 1000 MHz	2699.5			

Table 2. Run Times for the ZSU-23/4
and T-72

The values in Tables 1 and 2 are significant when deciding what kinds of resources are needed to produce low-frequency signatures. For the 1000 MHz case, memory requirements go through the roof while run times make it necessary to utilize multi-processor machines. The use of high-performance computers has provided the means by which these signatures can be produced.

ACKNOWLEDGEMENTS

REFERENCES

1. *Virtual Target Parametric Fidelity Analysis,* TMO Virtual Target Program, Document SED-VTL-TEC-005, SAIC, Huntsville, Alabama, March 1999.

2. Vechinski, Scott R., Deborah A. Potter, and James L. Cook, *Parametric Studies for SRBM RF Signatures: Solid Modeling Production Practices for SRBMs (U)*, SAIC, Huntsville, Alabama, August 1996.

3. James L. Cook, Clifford D. Geohagan, Samuel D. McKenzie, and Scott R. Vechinski, *Development & Validation of ATR Algorithms*, May 1998.

4. Vechinski, Scott R., and James L. Cook, *SAR Image Sensitivities to Virtual Target Detail and Resolution Variations (U)*, SAIC, Huntsville, Alabama, November 1996.

5. Schroeder, W., Martin, K., Lorensen, B., *The Visualization Toolkit*, Prentice Hall, Upper Saddle River, 1996.

6. Song, Jiming, FISC Developer, *private communication,* 1998.

7. *User's Manual for FISC (Fast Illinois Solver Code)*, Champaign, Illinois: DEMACO, Inc., 1997.

Joint Navy and Air Force Infrared Sensor Stimulator (IRSS) Program for Installed Systems Test Facilities (ISTFs)

Tom Joyner, NAWCAD/ACETEF, Patuxent River, MD/ COMPTEK-PRB Associates, Hollywood, MD
Bob Makar, Steve Jacobs, Dough McKee, Frank Stanley, COMPTEK-Amherst Systems, Inc., Buffalo, NY

ABSTRACT

The Office of the Secretary of Defense (OSD), Central Test and Evaluation Investment Program (CTEIP) is tasked to provide a coordinated process for making joint investments in defense test & evaluation (T&E) to offset the challenges presented by declining investments in test assets and increasing test requirements. Under CTEIP sponsorship, the Navy and Air Force are jointly developing three Joint Installed System Test Facility (JISTF) enhancements that are based on dynamic virtual reality simulation technology. The three enhancements are the Infrared Sensor Stimulator (IRSS), Generic Radar Target Generator (GRTG), and Joint Communications Simulator (JCS). The subject of this paper is the IRSS that was first briefed at the 1997 GTM&V conference.

The IRSS system will be used to stimulate *installed* Infrared/Ultraviolet (IR/UV) Electro-Optic (EO) sensors undergoing integrated developmental and operational testing. The IRSS generates digital infrared scenes in real-time to provide a realistic portrayal of infrared scene radiance as viewed by an IR system under test in a threat engagement scenario. This paper will describe the continuing IRSS development effort including new work completed in the past year. There will be a brief overview of the IRSS subsystems and functions, with emphasis on recent enhancements to its IR modeling capabilities. Specifically, the paper addresses issues involving the integration of three IR models: Spectral In-Band Radiance of Targets and Scenes (SPIRITS), Physically Reasonable IR Signature Model (PRISM), and IR Electro-optical Naval Engagement (IRENE). Also, there will be discussion regarding use of a radiometrically accurate method of employing geospecific material properties in the rendering of background terrain.

KEYWORDS: Installed Systems Testing, Infrared Sensors, Scene Simulation, Sensor Fusion, Interoperability, Electronic Combat Test Process, Infrared Scene Projection, Sensor Stimulation.

INTRODUCTION

The Infrared Sensor Stimulator (IRSS) is a modular cost-effective system that will be used to generate high fidelity Infrared (IR) scenes for stimulation of *installed* IR Electro-Optic (EO) sensors on aerospace platforms undergoing integrated developmental and operational testing. The IRSS will be capable of stimulating multiple types of sensors such as Forward looking Infrared (FLIR), Missile Warning Systems (MWS), Infrared Search and Track (IRST) and Missile Seekers. It is being developed under the sponsorship of the Office of the Secretary of Defense (OSD) Central Test and Evaluation Program (CTEIP) for use in a Joint Installed Systems Test Facility (JISTF) environment. The IRSS will be capable of satisfying installed sensor system test requirements through dynamic stimulation of IR/EO sensors which are integrated with other avionics processing software and platform sensor systems, [(e.g., radar, operational flight programs (OFP)]. To be a valid test tool, the spatial, spectral and temporal components of the IRSS computer-generated synthetic scenes must be of sufficient fidelity to produce sensor responses that are indistinguishable from the tested sensor's response to "real-world" conditions. This paper discusses the current capabilities and recent additions to the IRSS.

IRSS OVERVIEW

The IRSS System is an integrated hardware/software system that has been specifically designed to support the design, development, integration, and testing of IR/EO sensor systems. The IRSS supports both performance characterization and integrated sensor testing. The IRSS system generates radiometrically correct scenes in real-time for reactive installed systems testing of a variety of infrared and ultraviolet sensor systems. The generated scenes provide a realistic portrayal of the infrared scene radiance as viewed by the unit under test (UUT) in operational scenarios. Use of commercial-off-the-shelf (COTS) Silicon Graphics (SGI) fast symmetric multiprocessing hardware has minimized cost and development time. During real-time scene simulation, the multiprocessors are used to update polygon vertex locations and compute

radiometrically correct floating-point radiance values for each waveband. Scene radiance is calculated on a frame by frame basis accounting for the relevant contributions from the sky, sun, targets, terrain, and atmosphere as a function of the engagement geometry by using existing validated high-fidelity IR models.

The frame output of the IRSS system is configurable to match the characteristics of the sensor system under test. Sensor parameters such as frame size, frame rate, spectral band, number of bands, pixel resolution, and field of view are user configurable. The digital output of the IRSS can be formatted for direct injection into receiver/processor hardware or to drive an infrared projection system.

The baseline IRSS system includes the hardware and software components to provide a complete IR/EO simulation and test environment. Functionally, the IRSS system includes software to support offline modeling, database development, scenario generation,

and simulation control. Real-time functions include scene generation and sensor stimulation. The IRSS system supports both open-loop and closed-loop simulation. Open-loop simulation provides the user with the capability to execute predefined, time-sequenced scenarios ensuring total control over scenario events. Closed-loop simulation is supported through an external interface where the unit under test (UUT) and target position data can be generated by external simulations and provided to the IRSS system for reactive engagements.

In an integrated configuration, the IRSS can be coupled with Radio Frequency (RF) systems and facility-level composite mission simulators for correlated, synchronized multispectral testing. The IRSS supports the stimulation of single or multiple aperture sensor systems. The system is modular in design to support incremental expansion of both function and performance to meet current and future test requirements.
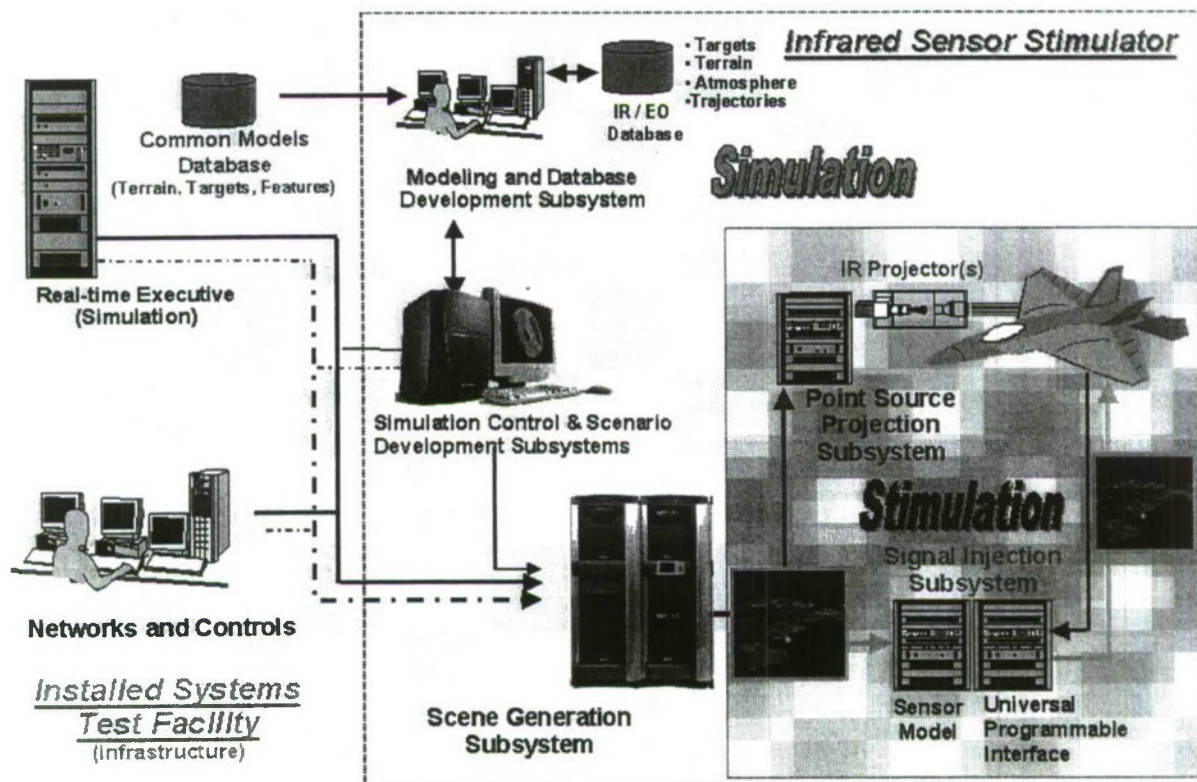
Figure 1 IRSS System Architecture

IRSS SYSTEM ARCHITECTURE

IRSS is a family of integrated software applications and hardware that supports all phases of the IR simulation and test process. Applications are available for offline modeling and scenario development, as well as real-time scene generation and sensor stimulation. The IRSS, as illustrated in figure 1, consists of six primary subsystems that are partitioned between offline and

real-time functions. The offline functions include the Modeling and Database Subsystem (MDBS) and Scenario Development Subsystem (SDS). These applications provide the user with all of the tools necessary to model and construct a virtual T&E warfare environment including terrain, targets, false targets, and atmospheric and weather parameters. The Simulation Control Subsystem (SCS) and the Scene Generation Subsystem (SGS) are the core of the system and provide the computing resources and processing required to generate infrared scenes in a real-time reactive mode. The two *stimulation* subsystems, the Signal Injection Subsystem (SIS) and the IR Point Source Projector (IRPSP) Subsystem, provide the capability for real-time electrical signal injection into the processing electronics and/or optical projection of scenes directly onto the sensor's detectors. COMPTEK-Amherst Systems Inc. (CASI) is developing the four subsystems of the scene generation/simulation component. COMPTEK-Amherst Systems Inc and SPARTA Inc. (SPARTA) are developing the SIS and IRPSP stimulation subsystems, respectively. A full field of view (FOV) image Scene Projection Subsystem (SPS) is planned as a future enhancement. Ultraviolet

generation and projection are also a planned future enhancement.

SCENE GENERATION CAPABILITY

The **MDBS** capability provides the test engineer or operator with the capability to build files representing threats, real and false targets, backgrounds, and atmospheric elements off-line. (e.g., In a non-real-time mode the operator will build files from sources such as plume radiance models, missile trajectory models, terrain elevation data, measured and/or statistically derived clutter data, and atmospheric models.) The primary output of the MDBS is the IR/EO Database, which contains the files used for subsequent scenario development and real-time simulation. The models identified in Table 1 are used in the calculation of signatures, atmospheric conditions, and target and test platform flight paths. Model selection is based upon degree of use in the simulation community, identified as a government 'standard', e.g. endorsement by Survivability Vulnerability Information Analysis Center (SURVIAC) or Joint Army, Navy, and Air Force/ Chemical Propulsion Information Agency (JANNAF/CPIA.

Table 1 – Third Party Model Utilization

Model	Function	Implementation
Signatures		
SPF/SIRRM	Missile & Air Vehicle Plumes	Point source intensity only
SIRRM	Extended Plumes	Under investigation
SPIRITS	Air Vehicle Body	,000 airframe facets typical esolution → facet reduction synchronous, non-real-time execution via interface
PRISM	Ground Vehicles Ships	,000-8,000 tank facets typical esolution → facet reduction synchronous, non-real-time execution via interface
TERTEM	Terrain Heat Transfer/ Temperature	OSART/ERTEM terrain thermal is integrated lack box interface for using other models
IRENE	Ships & Sea Backgrounds	hip signatures integrated as OpenFlight™ objects ea surface integrated as radiance textures
Atmospheric		
MODTRAN	IR Attenuation, Path Radiance & Solar Irradiance	eal-time lookup tables from offline execution ile based interface enables use of other models
OSIC	UV Background/Scattering	eal-time lookup tables from offline execution mplemented as prototype only
Cloud	Background (not 3D)	Under investigation
Obscurant	Background (not 3D)	Under investigation
Trajectory		
BLUEMAX	Test & Adversary Air Vehicle Flight Paths	Integrated or offline execution → scripted trajectory with interactive graphical way-point entry
ESAMS	Surface-to-Air (S/A) missile flyout	Integrated or offline execution → scripted trajectory
TRAP	Air-to-Air (A/A) missile flyout	Integrated or offline execution → scripted trajectory

The MDBS also supports importing and converting external database elements from common terrain or target databases that use a standardized open architecture, three-dimensional geometric file format to provide commonality with other ISTF stimulators. Extended OpenFlight™ has been selected as the 'standard' for model input/output and databases. The IRSS incorporates the MultiGen™ Application Program Interface (API) as a tool to support the creation, attribution, integration and execution of the models and databases. Use of the MultiGen also enables the import and manual attribution of other external databases. This process is illustrated schematically in Figure 2 using extended targets as an example. The construction of a Flight File Translator (FFT) is performed once for each external model that is to be used by the system. The primary objective of the FFT is to transform the model's native geometry representation into the OpenFlight format. The secondary objective of the FFT is to automatically place the appropriate object temperature and material attributes into the OpenFlight file. After this process is completed, the resulting OpenFlight files can be accessed and specified as scenario components through the IRSS Scenario Builder application. In the past year this process has been used to successfully import and utilize the outputs of PRISM, SPIRITS, and IRENE in IRSS scenarios.

Figure 2 – Target Model Integration

The output of the FFT is an OpenFlight file representative of the conditions for which the external model was executed. There are instances where some of the original conditions will change during the course of a scenario. Examples include tank barrel heating, engine compartment temperature, and throttle setting. These changes can be dynamically incorporated into a scenario through the Plug-In-Interface. This interface provides a mechanism by which specific changes can be incorporated into an object description when executing a scenario in real-time. The plug-in interface is a non-synchronized interface that enables third party or other external models to provide asynchronous updates to executing scenario files. The update frequency depends on model performance and the fidelity required for the target, and/or background signatures.

Model translators are interfaces to MDBS. They can be developed by CASI or by IRSS users. Each model translator can consist of a graphical user interface (GUI), a model processing function, and database translator. The GUI provides easy and efficient execution of the model. This feature is important when operators are unfamiliar with the specifics of each model. The model processing function compiles the model output and performs the necessary manipulation of data for real-time scene generation. The database translator formats the data into a common database format for scenario development and scene generation.

The MDBS contains the models, tools and databases used to represent targets and backgrounds in the test scenario. To compensate for their non-real-time execution speed, some models, e.g. Moderate Resolution Transmission (MODTRAN), Enhanced Surface-to-Air Missile Simulation (ESAMS), Trajectory Analysis Program (TRAP), are executed offline to create look-up tables or databases that are used during run-time scene generation. Once created, these look-up table databases become part of an EO/IR library. The IRSS System is required to respond to unscheduled, non-scripted events including man-in-the-loop commands in the external control-state whereas the trajectory models within the IRSS System are only intended for scripted applications. Not all of the models are restricted to off-line execution.

The **Scenario Development Subsystem (SDS)** provides the operator or test engineer with the capability to define simulation scenarios in which single or multi-sensor equipped vehicles move through a simulated test area. The output of the SDS is a scenario file that is saved in the scenario database. The scenario file references scripted terrain, targets, threats, trajectories, and special effects selected from the IR/EO database. The file also references customized simulation elements. These elements include an atmospheric specification, sensor specification, test platform assignment, sensor channel assignment(s), and state information such as situation display setup, visual display setup, and instrumentation setup. The SDS provides a convenient user interface for quickly building or editing scenarios based on libraries of objects created in the MDBS. An interactive situation display provides a graphic scenario building and display capability. The situation display features interactive control of viewing geometry, symbology, and scenario components. An interactive scenario sequencer provides the ability to setup scenario parameters and script scenario events. Used in conjunction with the situation display, the sequencer provides an efficient environment for building, editing, and previewing test scenarios. The IR models TRAP, Aircraft Flight Path Generator and Mission Performance Evaluation Model (BLUEMAX), and

ESAMS have recently been integrated with the SDS. As a result of this integration, the situation display now includes a scenario sequencer that can be used to provide an interactive graphics-based environment for the preparation of scripted trajectories. During the scenario development process, the scenario can be previewed using the scenario animator. This feature allows the operator to pre-run the test engagement and evaluate scenario events against the simulation timeline. The scenario gaming area, player motion, and UUT FOV are visualized in the situation display. Simulation clock controls are provided to stop, start, and pause the scenario.

The **Simulation Control Subsystem (SCS)** provides the operator or test engineer with the capability to control the execution of a simulation and perform fault tests on the IRSS channel hardware. A fault test and diagnostic capability is provided for assessing the health of the system and to assist in the operational

Integration of Extended Target Models with IRSS
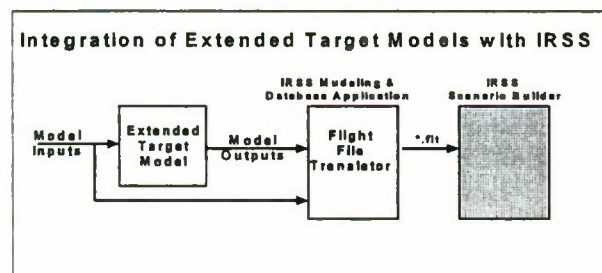
maintenance of system components. Input to the fault test function includes diagnostic scripts executed by the operator to determine the hardware operational state. Output consists of the pass/fail status of the performed tests along with status or trace messages showing test progress. Upon execution of the IRSS application, the SCS initializes by opening an existing or archived simulation file from the scenario database and setting the control state. When external and integrated control is disabled, the IRSS operates in a stand-alone mode in which the operator controls the simulation clock, situation display, visual display, and all instrumentation. When external control is enabled, control of the simulation clock, player positions and state, test platform interface, etc., is assumed by the Installed System Test Facility (ISTF) Operational Control Center (OCC). The OCC may send a load and initialize command to the SCS that contains a scenario script specifying some or all scenario and configuration information.

The **Scene Generation Subsystem (SGS)** produces IR/EO scenes in real-time. The term 'real-time' is relative to the frame rate of the sensor under test (e.g. 30 – 100 Hz for FLIR, 100 – 400 Hz for MWS). The SGS incorporates 'first principle' algorithms for the

radiometric signature computation. A 'virtual' test may involve stimulation of up to three sensors requiring multiple SGS channels. Each channel stimulates a single sensor or a single aperture of a multi-aperture sensor. A sensor-specific configuration is supplied during initialization. The SGS performs both scene generation and scene rendering. During scene generation, the SGS determines the test scenario viewed gaming area, on a frame to frame basis, based on the direction or view of the sensor line-of-sight, and host platform position in space (e.g. altitude, heading, pitch, roll). The specified simulation file is examined to determine which polygons, representing players and targets, and background elements, occur within the viewed area, and are to be displayed in the simulation. The material characteristics and polygon viewing geometry are used to calculate radiometric values for polygon vertices. During scene rendering, polygons are decomposed into pixel elements and inserted into an output frame buffer resulting in a radiometric, spatial, and temporal representation of the scene as viewed by the sensor relative to its line-of-sight. This digital scene is the input from the SGS into either the SIS, for conversion into an electrical signal that is injected into the sensor processing electronics, or the IRPSP, for optical projection into the sensor's entrance aperture. Additional discussion of this subsystem is presented below.

The **IRSS Scene Generator** generates radiometrically accurate scenes for installed systems testing of a wide variety of IR and UV sensor systems. The generated scenes provide a realistic portrayal of the in-band scene radiance as viewed by the system under test in operational scenarios. Scene radiance is calculated on a frame by frame basis accounting for the relevant contributions from the sky, sun, targets, terrain, and atmosphere as a function of the engagement geometry. The frame output of the scene generator is configurable to match the characteristics of the sensor system under test. Sensor parameters such as frame size, frame rate, spectral band, number of bands, pixel resolution, and FOV are user configurable. The digital output can be formatted for direct injection into receiver and processor hardware or to drive an IR projection system.

The IRSS Scene Generator was designed specifically to address the core technical issues for IR/EO scene generation. Commercial scene generation systems are optimized for visual effects and standard display devices. Real-time IR/EO sensor stimulation requires a higher level of fidelity (scene quality and radiometric content) and usually involves large frame sizes at high

frame rates. The current scene generator hardware configuration consists of a Silicon Graphics Onyx2 InfiniteReality® graphics computer with eight or twelve 'R10000' processors. The system can use either the SGI InfiniteReality or the COMPTEK-Amherst Systems Scene Rendering Subsystem (SRS) for final image rendering. The SRS is designed specifically for infrared applications while the InfiniteReality is optimized for visual applications. The selection of the graphic system depends on the objective of the test facility. When evaluating the detection, tracking, or guidance performance of a sensor system, fidelity and radiometric validity are critical. In this situation, the accuracy and programmability of the SRS may be required. In cases where radiometric accuracy is less important and validation is not an issue, the InfiniteReality option may be preferred.

The **IRSS Scene Rendering Subsystem (SRS)** is a graphics-processing pipeline developed specifically for rendering IR/EO scenes. The SRS overcomes many of the problems associated with adapting visual rendering systems for IR/EO simulation. The SRS uses full 32 bit floating point accuracy for all calculations including radiance (lighting), transparency, texture mapping and filtering, anti-aliasing, and hidden surface removal (z-buffering). The SRS can process up to six 16 bit colors or three 32 bit colors per pixel. Equations for pixel level lighting and atmosphere effects can be modified as desired to make tradeoffs between rendering accuracy and speed. Depending on the tradeoffs selected, a fully configured SRS can provide more than four times the radiometric accuracy of the InfiniteReality. In addition, the SRS can be tightly coupled with an IR/EO sensor in a reactive, closed-loop configuration, and dynamic frame size and frame rate changes can be processed with low latency.

SCENE PRESENTATION CAPABILITY

The **Signal Injection Subsystem (SIS)** accepts digital scenes produced by the SGS and creates an electrical digital or analog signal that is injected into the sensor image and signal processing chain. This subsystem is currently being developed and will be manufactured by CASI as a deliverable under an Air Force SBIR contract for a Universal Programmable Interface (UPI).

As part of the signal creation function, the SIS must modify the scene to represent the effects of bypassed sensor components and phenomenology prior to the injection point, convert the modified image to a properly conditioned electrical signal, and provide the electrical connection to the sensor. The scene

modification is accomplished by two custom processing components within the SIS: a convolution processor and a pixel processor. These SIS components use digital signal processor (DSP) arrays; high-speed Xilinx programmable gate array chips for the convolution processor, and Motorola 266 MHz Power PC 740 chips for the pixel processor. The latter hardware assembly is common to the polygon processor in the previously discussed CASI rendering engine. The sensor interface module (SIM), which is unique to each sensor, provides generation and conditioning of the electrical signal and its physical connection to the sensor. This assembly is a plug-in module that enables the SIS to be easily configured for different sensors.

Additional functions of the SIS include processing (e.g. I/O handling) sensor control signals and, if necessary, emulating their functionality. These functions may require one or more electrical connections to the sensor or other test platform avionics systems. The SIS is based on the CASI Universal Programmable Interface UPI, which is discussed in more detail later in this paper.

The **IR Point Source Projector (IRPSP)** is another stimulator sub-subsystem that presents a generated scene to the sensor. The primary function of the IRPSP is to accept digital input scenes produced by the SGS and to generate equivalent output scenes, in the form of in-band EO/IR energy, for projection into the entrance aperture of the UUT. The format of the scene input to the IRPSP from the SGS will be SGI Direct Digital Output for the Onyx2 (DDO2), also known as the Onyx2 Digital Video Port (DVP). The IRPSP will also be capable of receiving scene input from the SIS in a DDO2 format. Setup and control of the IRPSP will be managed by the SCS via the SGS to IRPSP interface. The IRPSP will consist of seven primary subsystems. These subsystems are the Control Electronics Subsystem (CES), Environment Control Subsystem (ECS), Infrared Emitter Subsystem (IRES), Mounting Platform Subsystem (MPS), Non-Uniformity Correction Subsystem (NUCS), Projection Optics Subsystem (POS), and Software Control Subsystem (SCS).

NEW CAPABILITIES

Advances in Maritime Modeling

A requirement for IRSS to test the U.S. Navy's AAS-44V FLIR system led to the introduction of a Maritime Combat Environment (MACE) modeling capability into IRSS. This involved the integration of a maritime thermal model derived from the U.S. Navy's IRENE model. The fact that the integration of this model was a smooth process is due to two main factors: cooperation between the developers of Infrared Electro-optical Engagement Model (IRENE) and CASI and the easily accessible OpenFlight format used by IRSS. The cooperation between the two parties allowed the work to be done in a minimal amount of time. Also, the expertise of the U.S. Navy development team allowed for great control of the way in which IRENE could be used for implementation. The API component of the MultiGen-Paradigm's Creator™ program enables easy access to the OpenFlight format. Consequently, conversion of the IRENE file format to the IRSS-supported OpenFlight format was very straightforward.

A major portion of the MACE effort involved the development of a method for the creation and rendering of the ocean background. The background is generated by a ray-tracing routine based on The Naval Research Laboratory's Kelvin and Random Ambient Sea Waves (KELSEA) model, which computes the source radiance of each square texel in a 512 x 512 grid. These texels can have a size of 1m or 5m on an edge, resulting in higher or lower resolution-radiance map textures. These radiance maps are then rendered by IRSS, which computes the atmospheric effects.

The MACE team is currently seeking to identify sources of future funding for the continued development of the MACE capability. Some of the features earmarked for future work include the rendering of wakes, the creation of sea height maps for ocean backgrounds, and the inclusion of plumes in ship models. Figure 3 is a sample maritime image generated by IRSS.

Advances in Terrain Simulation

Terrain definitions are fully attributed faceted surface descriptions derived from Digital Terrain Elevation Data (DTED) augmented with cultural details such as roads, bridges, and buildings. The DTED data is used to create polygonal wire-frames representing terrain contour or shape. Terrain attributions include material properties, textures, and temperature specifications. Background detail (e.g., texture) at the sensor pixel level is represented by texture maps overlaid on larger terrain polygons.
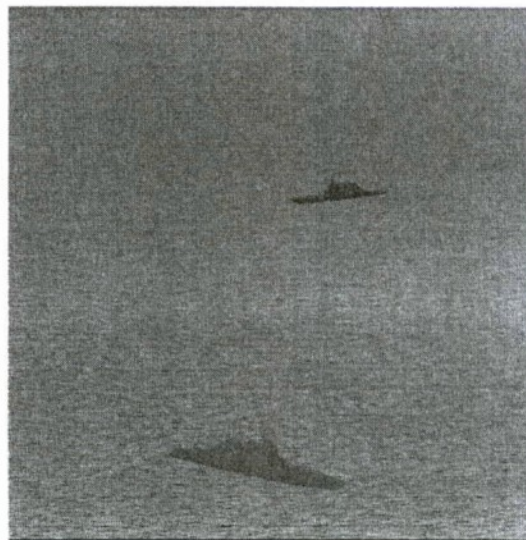
Figure 3 Example Maritime Scene
(The ocean surface texture is 1m resolution and a sea state of 2. The sensor altitude is ~1400m with a look angle of −10°. Ranges to the two ships are 740m and 1560m. White is hot.)

Radiometrically-correct real-time simulation of realistic terrain requires three essential elements. First, a high resolution description of the physical properties of the terrain, both in terms of material composition and topography is needed. Second, high-fidelity models for the sensor and its physical environment must be employed. Finally, sophisticated algorithms must be employed to combine the models and the terrain description into rendered scenes accurately and in real-time.

The IRSS effectively combines these three elements, providing a new level of realism to real-time sensor simulation. Since the terrain description is based only on its physical properties, it can be used to simulate the terrain regardless of the waveband(s) of the sensor being modeled, and correlation of different waveband images is easily accomplished. The description can also be used in conjunction with a thermal model to include realistic seasonal and diurnal effects. Radiometric accuracy is achieved through the use of accepted phenomenological models and advanced algorithms. Geospecific texturing results when correlated satellite imagery and digital elevation models for a specific region are used to create the terrain description.

Terrain Description and the Models

Increased availability of satellite imagery, and the development of sophisticated image analysis techniques, has made the high-resolution description of terrain material composition a practical reality. Classification techniques are employed to determine the material or material mix of the terrain from satellite images on a texel-by-texel basis. A material code number is assigned to each texel, and all the codes for a specific patch of terrain are assembled into a 'material map', or, in the case of a material *mixture* being assigned to a texel, a 'material mix map'. The material codes are cross-referenced to a table that gives the pertinent properties of each material. The use of material mixtures has an advantage over using a single material per texel in that it enhances the level of detail in the terrain image and smoothes the transitions between regions of differing material types.

Two types of topographic descriptions of the terrain are required. First, the effective utilization of computer graphics technology drives the need for the terrain to be described in terms of a triangular irregular network (TIN). A TIN representation is readily obtained from government-distributed digital elevation data by using commercially available Delaunay algorithms.

The second type of topographic description required is at a higher, texel-level, resolution. This is necessary due to the sensitivity of the texel's radiance to its normal vector and its elevation. While the texel's source radiance is modeled as being independent of its orientation (i.e., Lambertian), its normal vector and elevation can have a significant impact on the source radiance, by effecting the texel temperature. The normal vector also determines how much sunshine, skyshine, and earthshine the texel reflects. The texel-level topographic data is readily derived from digital elevation data using standard interpolation and gradient estimation techniques.

To efficiently store texel-level terrain data a new file format, material mix and topography (MMT), was developed for IRSS. This format stores the material mix data for each texel, as well as the texel-level topographic data, into a single file, which is then correlated to a TIN in the same manner as a normal texture. The topographic data takes the form of elevation, 2-D gradient, and cross-derivative samples at equally spaced posts. Elevation and normal vector data is then easily calculated at intermediate texel locations using bicubic interpolation. This bicubic representation itself reduces the storage requirement from 16 bytes/texel to less than about 1 byte/texel. The post spacing is selected to approximate the resolution of the

source data, which can result in further efficiency without adding any additional processing burden.

The primary models employed by IRSS in terrain simulation are for sensor spectral response, atmospheric effects, and the determination of terrain temperatures. The user models the sensor spectral response during the scenario development process by simply entering sensor response values, and corresponding wavelengths, into a table. Atmospheric effects are modeled using the industry-standard atmosphere model, MODTRAN. The IRSS architecture is designed to facilitate the use of different thermal models, but currently uses only Terrain Temperature Model (TERTEM).
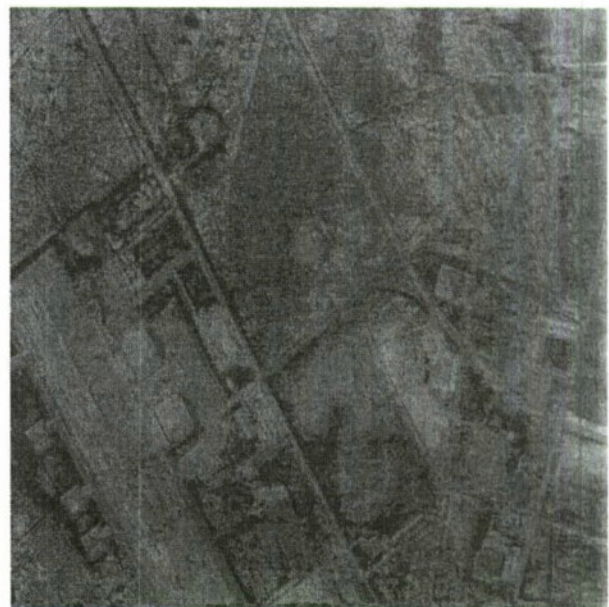
Figure 4 Radiance Map Terrain

Terrain-rendering Algorithms

For maximum efficiency, the terrain-rendering algorithms are carefully designed to perform optimized pre-run-time calculations while also preserving accuracy. This non-real-time calculation consists of the generation of lookup tables and texture. The lookup tables are used to calculate attributes for terrain facet vertices that vary widely with the position of the sensor relative to points on the terrain. These attributes account for the effects of atmospheric attenuation and path radiance. The texture is used to account for a number of first principle physical effects including temperature variations, thermal radiation, and solar, skyshine, and earthshine reflections. The generation of this texture, called an adjusted radiance map, is expedited by first generating, and then using, lookup tables. Once the lookup tables, and adjusted radiance

map, are precalculated, the real-time portion of the simulation can begin. The attributes of terrain facets within the field-of-view are calculated on a vertex-by-vertex basis, and the results sent to a rendering engine along with the specially-formulated texture. The IRSS has the capability to render scenes using either SGI graphics hardware, such as the InfiniteReality, or by using CASI's SRS, which is designed specifically for sensor applications. In either case, a unique rendering algorithm is employed to create the desired imagery with high accuracy.

and gimbals must be emulated. Real-time sensor modeling must be performed to correctly model the by-passed sensor optics and electronics for the case of direct signal injection. Additionally, optical projection requires non-uniformity correction (NUC) of the thermal array.

A Universal Programmable Interface (UPI) has been developed under the IRSS program to provide such functionality. Unlike custom solutions, the UPI provides a reconfigurable method for interfacing a wide range of UUTs through either direct injection or optical
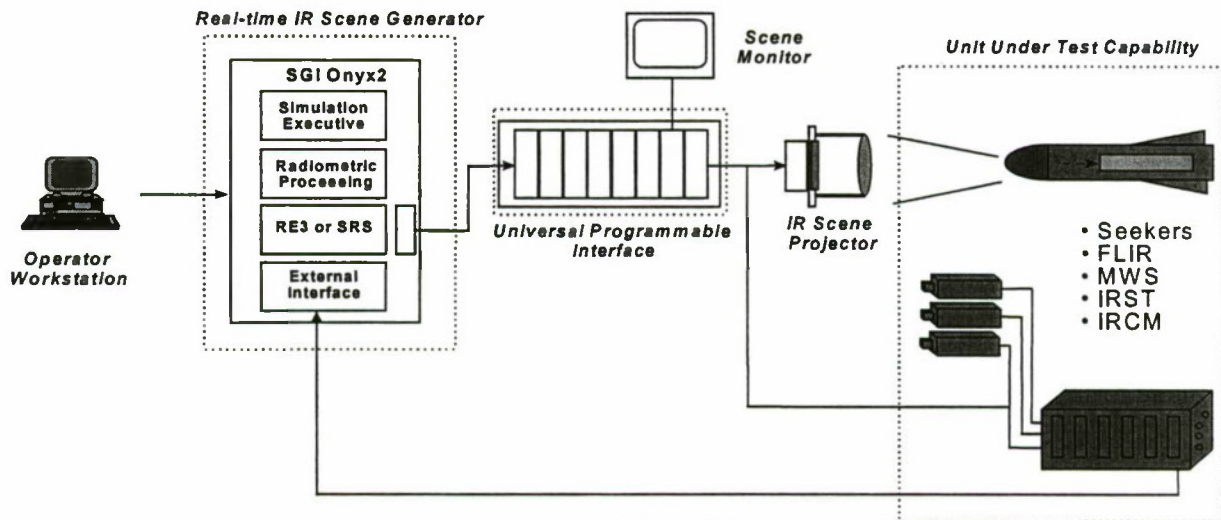
Figure 5 IR/EO HWIL Testing with a UPI

To assess the accuracy of this new rendering technique, precise calculations of the apparent radiance of the terrain were compared to the results that would be obtained using the rendering technique, for a wide variety of sensor-to-terrain geometries and parameter variations. This showed that the error introduced by the algorithms used were generally a fraction of a percent, but that in certain extreme cases can grow to approximately 1%.

Advances in Sensor Simulation

IR/EO sensor system testing requires valid stimulation of the UUT to correctly determine the performance of the sensor system and processing algorithms. Assuming the scene generation system (SGS) has correctly modeled the target and background signature and atmospheric attenuation, other requirements exist for valid stimulation. A physical interface is required between the SGS and the UUT to properly reformat the data such that it can be introduced to the UUT either through the direct injection of the signals into the system's processing electronics or the projection of in-band scene radiance into the sensor's optical aperture. By-passed missing components such as gyros

projection. This flexible capability is achieved through the use of a core architecture that provides industry standard interfaces, coupled with minimal custom interface hardware and reconfigurable software. The UPI provides the physical interface between an SGS and the UUT, performs sensor modeling, and emulates missing components. An illustration of using the UPI in hardware-in-the-loop (HWIL) sensor system testing is shown in Figure 5.

Both current and notional sensors can be modeled with the UPI, providing system designers the capability to measure the effect on performance due to changes in sensor design. Additionally, the UPI can be used in applications requiring high speed general purpose image processing capabilities.

The top-level architecture of the UPI is illustrated in Figure 6.

```
                                    ┌──────────┐
                                    │   CRT    │
                                    │ Display  │
                                    └──────────┘
                                         ↑
From IR/EO       ┌───────────┐  ┌───────────┐  ┌───────────┐      To IR/EO
 Scene     ────> │ Convolver │=>│   Pixel   │=>│  Output   │───>   System
Generator        │           │  │ Processor │  │Controller │
                 └───────────┘  └───────────┘  └───────────┘
```
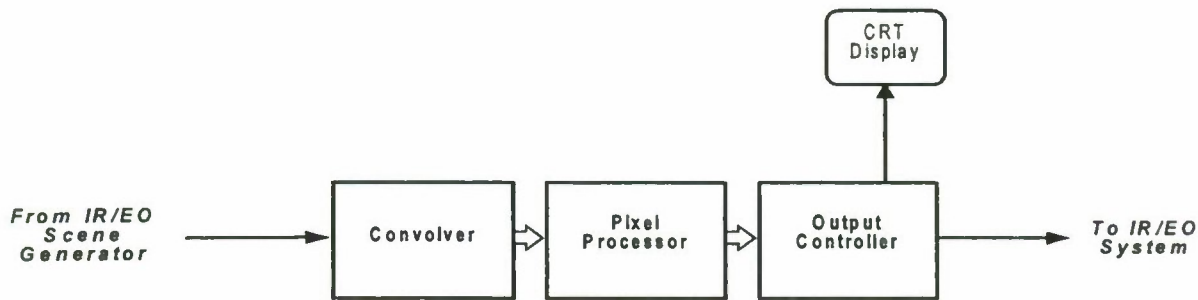
upi-033-032899

Figure 6 UPI Top-level Architecture

Convolver

For direct signal injection, the UPI must accurately model the modulation transfer function (MTF) effects of the bypassed sensor optics and electronics. The MTF is modeled through the convolution of a spatial kernel in the convolution subsystem (convolver). This kernel is generated offline through the use of the UPI sensor modeling software developed at CASI. An industry standard tool for IR sensor characterization and performance modeling (FLIR92), was chosen as a basis for this software. Based on the input sensor parameters, this tool provides the user with the ability to generate the convolution kernels that model the sensor MTF. The user can model the MTF of current and future sensors

In the convolution subsystem, the over-sampled rendered scene is convolved with the MTF kernel, producing the image as seen by the sensor. Ideally, a static mapping between the sensor pixels and the scene pixels would exist. However, multiple factors can contribute to sensor pixel displacement from the perfectly rendered image to the correctly sensed image, including scene-rendering latency, optics-induced geometric distortion, and physical sensor jitter. In addition to convolution, the convolution subsystem handles these factors through displacement processing.

In a closed-loop installed-systems test configuration, latency can occur between the time when positional data is received and the scene is rendered. Latency can create errors in x and y shift, and in rotation, which can consequently affect the ability to accurately test the system under test (SUT). This is a problem for both direct signal injection and projection. Therefore, the UPI performs latency compensation to extract the correctly located sensor image from within the oversized scene image.

Within a real sensor, the optics generate geometric distortions in the sensed scene. Edges that are geometrically straight appear curved, an effect that is especially pronounced with wide FOV optics. Accurate emulation of the optics-induced distortion may be required when performing direct signal injection since the optics are bypassed. If the SGS does not perform the geometric distortion, then the compensation can be performed by the UPI as an additional sensor effect modeled by the system.

Sensors can experience physical jitter when mounted on a moving platform. For a scanning system, this could create small perturbations in the location of each successive scan line. If jitter affects the system performance, it should be accurately modeled. Assuming the jitter function can be mathematically modeled, it too can be handled by the UPI. Additionally, a variety of other user-defined mathematical displacement effects can be modeled.

Pixel Processor

A variety of noise sources and response non-linearity's that are present in sensors can affect performance. When performing direct signal injection, the sensor is bypassed. Therefore, these pixel effects are modeled, based on user specified sensor parameters, in the pixel processor to add to the validity of the installed systems test.

Some of the pixel effects that have been modeled include conversion from radiance to photons, various noise sources, linear responsiveness and automatic gain control (AGC).

The pixel processor utilizes a COTS circuit card assembly that contains general purpose processors. This board was originally designed for and used in the CASI SRS. Since the pixel effects are implemented as software executing on the general-purpose processors in the pixel processor, a wide variety of effects can be modeled in the UPI. User-defined effects, within the limits of the UPI, can easily be added with no change or cost in hardware.

Optical projection can also benefit from the pixel processor. Gain and offset tables can be loaded in the resident memory and used for NUC.

CONCLUSIONS

The Infrared Sensor Stimulator (IRSS) is a cost-effective system that provides flexible, re-configurable, reproducible, and repeatable full test environments for evaluating Electro-optical/Infrared (EO/IR) sensor systems during the concept, research, development, prototype, and test and evaluation phases. When employed as an integrated Installed System Test Facility (ISTF) element, it is a valid and verifiable test and evaluation risk reduction tool that optimizes use of costly range testing. The IRSS sensor modeling capability contributes to the development of systems and sensor, and engineering model development (EMD) performance effectiveness evaluation.

The scene-generation component has successfully completed its Spiral 4 development and the software has been delivered to the Navy and Air Force for evaluation. The signal injection component has completed preliminary design review, and the infrared point source projection component has completed the critical design review for both systems. Final IRSS/facility integration is scheduled for third. Quarter Government Fiscal Year (GFY) 2000, and will occur at the Air Combat Environment Test and Evaluation Facility (ACETEF), Naval Air Warfare Center – Aircraft Division (NAWC-AD), Patuxent River, Maryland and the Avionics Test and Integration Complex (ATIC), Air Force Flight Test Center (AFFTC), Edwards AFB, California. The system's Fully Operational Capability (FOC) completion is scheduled for the fourth quarter of GFY 2000.

REFERENCES

[1] O'Toole, B. E., "Real-time Infrared Scene Simulator (RISS)," Robert Lee Murrer, Jr., Ed., *Proceedings of SPIE*, **2741**, 209-218, Orlando, FL, 1996.

[2] Joyner, T., "Real-time IR/UV EO Sensor Stimulation," *Proceedings of the Ninth Annual Ground Target Modeling & Validation Conference*, 166-180, Houghton, MI, 1998.

Real-Time Sensor Modeling Using A Universal Programmable Interface

Douglas C. McKee

Comptek Amherst Systems, Inc

Buffalo, NY 14221

ABSTRACT

Infrared Electro-optical (IR/EO) sensor system testing requires valid stimulation of the unit under test (UUT). Prerequisite is a scene generation system (SGS) that produces imagery that correctly models the target/background signature and atmospheric attenuation. A physical interface is required between the SGS and the sensor UUT to properly reformat the data such that it can be introduced to the UUT either through the direct injection of the signals into the system's processing electronics, or the projection of in-band scene radiance into the sensor's optical aperture. A Universal Programmable Interface (UPI) for signal injection to the sensor is being developed by Comptek Amherst Systems, Inc (CASI) under contract to the Air Force Flight Test Center (AFFTC) at Edwards AFB, with additional funding support from the Naval Air Warfare Center – Aircraft Division (NAWC/AD) at Patuxent River, Maryland. Additionally, the UPI will perform accurate, real-time sensor modeling, which is a key element in the image signal injection pipeline. Unlike custom interfaces, the UPI is reconfigurable, providing the capability to interface to a wide range of sensor systems.

This paper will discuss the sensor modeling capabilities of the UPI and the supporting software and hardware. Sensor modeling capabilities include image blurring associated with the sensor's modulation transfer function (MTF) and pixel effects. A sensor modeling and analysis software tool, based on FLIR92[1], will be discussed. A technique for accurately performing latency compensation and modeling geometric distortion, physical sensor jitter, and other user-specified effects will also be discussed.

Keywords: Sensor Modeling, Infrared, Electro-Optic, Simulation

1. INTRODUCTION

The IR/EO sensor system testing requires valid stimulation of the UUT to correctly determine the performance of the sensor system and processing algorithms. In addition to correctly modeled target/background signature and atmospheric attenuation, other requirements exist for valid stimulation. A physical interface is required between the SGS and the UUT to properly reformat the data so that it can be introduced to the UUT either through the direct injection of the signals into the system's processing electronics, or the projection of in-band scene radiance into the sensor's optical aperture. By-passed/missing components such as gyros and gimbals must be emulated. Real-time sensor modeling must be performed to correctly model the by-passed sensor optics and electronics for the case of direct signal injection. Additionally, optical projection requires non-uniformity correction (NUC) of the thermal array.

A UPI is being developed to provide such functionality. Unlike custom solutions, the UPI provides a reconfigurable method for interfacing a wide range of UUTs through either direct injection or optical projection. This flexible capability is achieved through the use of a core architecture that provides industry standard interfaces, coupled with reconfigurable software and minimal custom interface hardware. The UPI provides the physical interface between an SGS and the UUT, performs sensor modeling, and emulates missing components. Both existing and notional sensors can be modeled with the UPI, providing system designers the capability to measure the effect of possible sensor characteristic design changes on performance. Additionally, the UPI can be used in applications requiring high-speed general purpose image processing capabilities.

The top-level architecture of the UPI is illustrated in Figure 1-1.
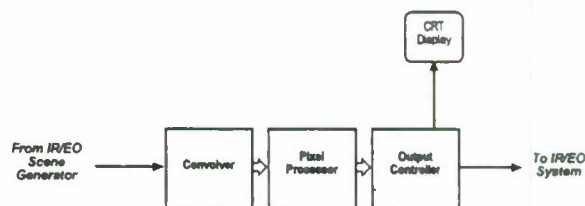
Figure 1-1. UPI Top-level Architecture

The input scene is blurred in the convolver. In parallel, the convolver performs latency compensation. The pixel processor applies pixel effects, including noise and linear responsivity. Additionally, the pixel processor provides NUC capabilities for optical projection.

Hardware and sensor requirements have been presented in the past.[2, 3] This paper focuses on the sensor modeling capabilities supported by the UPI, both in software and in hardware. Section 2 addresses the software used to model image blurring due to the sensor's MTF, and pixel effects. A sensor modeling and analysis software tool based on FLIR92 is used to generate the convolution kernel and simulate the pixel effects. Section 3 provides more detail on MTF modeling. Section 4 discusses pixel displacement processing, a technique that is used to accurately perform latency compensation and to model other sensor effects, including geometric distortion, physical sensor jitter, and various additional user-defined effects. Section 5 discusses pixel effects modeling. Conclusions are presented in Section 6.

2. SENSOR MODELING SOFTWARE

A key capability requirement for the UPI is to accurately model the sensor UUT characteristics, including MTF and pixel effects. The sensor MTF is the frequency representation that describes image blurring. Pixel effects include various noise sources and pixel nonuniformity. These effects are emulated in hardware components of the UPI; the convolution and pixel processing subsystems, respectively. Off-line software tools are necessary to build the required inputs for these subsystems.

Software tools are currently being developed by Comptek Amherst Systems, Inc. (CAS) for sensor modeling. FLIR92, an industry standard tool for IR sensor characterization and performance modeling, was chosen as a basis for this software. A graphical user interface (GUI) has been developed that provides the capability for a user to enter sensor parameters matching those used as input by FLIR92. The UPI sensor modeling tool also supports both the import and export of FLIR92 files to provide file compatibility with this model. After the appropriate parameters are specified, this tool provides the user with the ability to generate the convolution kernels that model the sensor MTF. The details behind kernel generation are provided later in this paper.

Software-generated sensor modeling data files add flexibility to the system. The tool's present functional capabilities provide the user with the ability to regenerate the kernels with modifications to the sensor parameters. Therefore, the user can easily model numerous existing sensors as well as potential future designs with no cost in hardware. Multiple sets of convolution kernels can be generated and stored for a given sensor system. This capability supports emulation of a multiple field-of-view (FOV) FLIR. When the FLIR's FOV is changed during a simulation run, the convolution subsystem will load in the set of kernels generated for the new FOV. Slight variations in each sensor's MTF in a multiple aperture system can also be emulated by generating a specific kernel for each sensor. Unique kernel sets can be generated for each sensor with small changes in the sensor parameters. Although FLIR92 is used as a baseline for the UPI sensor modeling tool, other more recent models such as NV-THERM can be incorporated with FLIR92 in the future. Due to the malleability of software, future sensor modeling refinements can be implemented with no additional hardware cost.

3. MTF MODELING

For direct signal injection, the UPI must accurately model the MTF effects of the by-passed sensor optics and electronics. The MTF is modeled through the convolution of a spatial kernel in the convolution subsystem. This kernel is generated off-line through the use of the UPI sensor modeling software. Similar to FLIR92, the overall system MTF is modeled as the combination of component MTFs. Components include diffraction MTF, geometric blur MTF, detector spatial MTF, detector temporal MTF, Focal Plane Array (FPA) integration MTF electronics low-pass MTF, etc. Each component's MTF is mathematically generated based on the sensor parameters, using the same equations as FLIR92. However, an exception does exist with the electronics low-pass MTF. FLIR92 only provides the magnitude of this MTF, but phase information is also needed to build a convolution kernel. Since the component MTFs are built in the frequency domain, combination is performed by multiplying them together. The software will additionally permit the input of measured MTF components in either the frequency or the spatial domain. If a component is specified in the spatial domain, it will be converted to the frequency domain through a Fast Fourier Transform (FFT).

The convolution subsystem utilizes a spatial domain convolution kernel. Therefore, after the MTF is generated, it is converted to the spatial domain through an inverse FFT. The MTF is generated using a large array so that it can be accurately represented across multiple frequencies. This approach results in a large convolution kernel. However, practical UPI implementation constraints dictate that the kernel size be limited to 16 × 16 due to computational cost. Therefore, the kernel is truncated to the desired size and quantized to the appropriate fixed-point resolution.

Although truncating the kernel lowers the computational cost, it can lead to inaccuracies and aliasing. The UPI sensor modeling tool provides the capability for the user to vary the kernel size and display the effects on the modeled MTF of truncation. An example comparison is displayed in Figure 3-1. The first image represents the 2D MTF prior to truncation of the kernel. The post-truncation MTF was created by transforming the truncated kernel back into the frequency domain through an FFT. Clearly, the truncated kernel does not perfectly reproduce the MTF modeled with a nontruncated kernel. The ringing at higher frequencies can lead to aliasing distortions. However, most of the ringing is along the horizontal and vertical axes with only minor ringing along the diagonals. It should be mentioned that these images were raised to the 0.25 power for illustrative purposes; the ringing isn't as pronounced as displayed.

Figure 3-1. UPI MTF. a) Pre-truncation,

Figure 3-1. UPI MTF. b) Post-truncation, raised to 0.25

The user also has the option to display the MTF absolute differences, and/or use log scales to fully analyze the differences due to truncation. Although differences will exist, this tool can be used to help determine the best tradeoff between kernel size and error. Convolution is performed on an over-sampled input scene. Therefore, the over-sampling factor can additionally be adjusted and its impact can be analyzed.

4. PIXEL DISPLACEMENT PROCESSING

In the convolution subsystem, the over-sampled rendered scene is convolved with the MTF kernel, producing the image as seen by the sensor. Ideally, a static mapping between the sensor pixels and the scene pixels would exist. However, multiple factors can contribute to sensor pixel displacement from the perfectly rendered image to the correctly sensed image, including scene rendering latency, optics induced geometric distortion, and physical sensor jitter. The following discussion describes how these effects can be modeled in the UPI through pixel displacement processing. Accuracy concerns will be addressed including a method for subscene pixel mapping.

Displacement processing requires a mapping from each sensor pixel into scene space. Each map entry contains an x, y pair specifying where this sensor pixel is centered, in scene coordinates. During convolution, the kernel is applied to the location in the scene based on the map entry for the current sensor pixel.

An initial map must be generated which is subsequently modified by displacement processing. This map can either be stored explicitly or expressed as a mathematical function. An example map is shown in Figure 4-1 in which the sensor pixels are mapped into the center of the generated scene. The first image is the input scene image with the overlaid squares representing a subsampling of the displacement map entries. The area covered by the sensor is less than that covered by the scene to allow for the displacement processing. The second image represents the output scene image produced by applying the convolution kernel at entries represented by the displacement map. If there were no pixel displacement, the initial map would be used directly for convolution processing. However, modifications to the map are necessary to model the above-mentioned effects.

Figure 4-1. b) Initial Sensor to Scene Pixel Displacement Map

4.1 Latency Compensation

In a hardware-in-the-loop (HWIL) system, latency can occur between the time when positional data is received and when the scene is rendered. Latency can create errors in x and y shift, and in rotation, which can consequently affect the ability to accurately test the system under test (SUT). Therefore, the UPI performs latency compensation to extract the correctly located sensor image from within the oversized scene image. This can be accomplished by mathematically rotating and shifting the x, y pairs in the displacement map. Rotation is performed through the following equations:

$$x' = x \cos(\theta) - y \sin(\theta)$$
$$y' = x \sin(\theta) + y \cos(\theta)$$

(1)

where x, y are the current coordinates stored in the displacement map; x', y' are the new coordinates and theta is the angle of rotation. Care must be taken to ensure that the rotation takes place around the boresight of the sensor. The implementation may require a shift operation to accommodate this requirement.

Image shift is calculated as:

$$x' = x + d_x$$
$$y' = y + d_y$$

(2)

where x, y are the current coordinates stored in the displacement map; x', y' are the new coordinate; and d_x, d_y

Figure 4-1. a) Initial Sensor to Scene Pixel Displacement Map

are the shifts in x and y, respectively. Any postrotation shift can be combined with this shift.

Figure 4-2 shows the result of rotation and shift of the displacement map coordinates. It should be noted that the clockwise rotation of the displacement map yields a
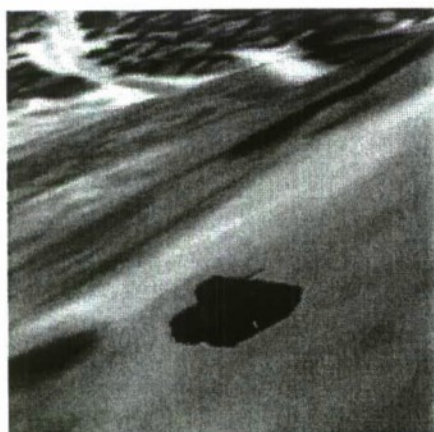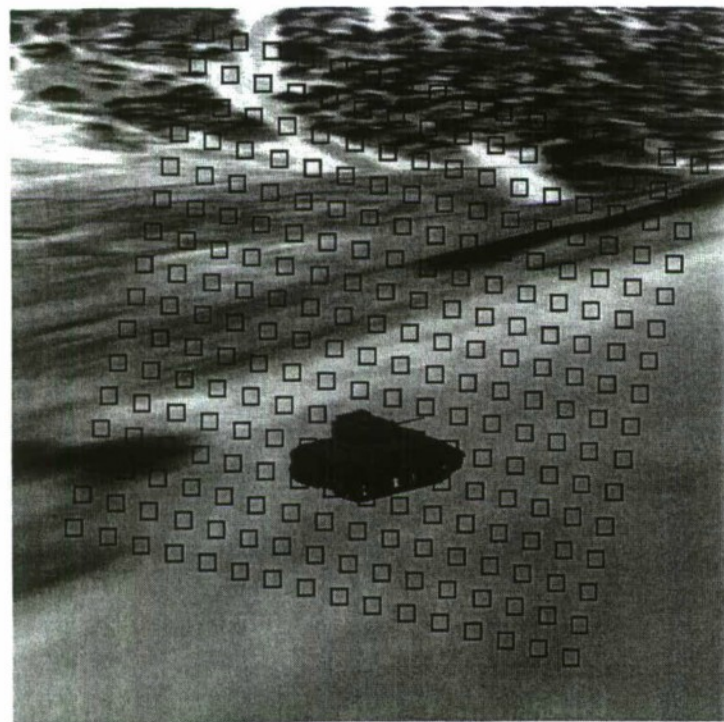

Figure 4-2. Rotated and Shifted Displacement Map

counter-clockwise rotation because the displacement map defines a mapping from the sensor to the scene, i.e., this is where the sensor image is extracted from the scene. Similarly, shifts of the displacement map lead to shifts in the opposite direction of the resultant sensor image.

4.2 Geometric Distortion

Within a real sensor, the optics can generate geometric distortions in the sensed scene. Edges that are geometrically straight can appear curved, an effect that can be especially pronounced with wide FOV optics. Accurate emulation of the optics-induced distortion may be required when performing direct signal injection since the optics are by-passed. If the scene generation system (SGS) does not perform the geometric distortion, then the compensation can be performed by the UPI as an additional sensor effect modeled by the system.

Two common geometric distortions are pincushion and barrel effects. In both cases, there is a mapping between the true off-axis distance and the imaged off-axis distance. A function that generates these effects in polar coordinates is:

$$\rho' = \left(k\rho\right)^{\alpha}\rho \tag{3}$$

where ρ is the off-axis distance; k is a scaling factor; α is the distortion exponent; and ρ' is the resultant off-axis distance. The distance, ρ, is multiplied by the scaling factor, k: setting k equal to one over the maximum distance normalizes this factor to between 0 and 1. The distortion exponent, α, determines both the type and the amount of distortion. A positive α generates a pincushioning of the displacement map, while a negative α generates a barreling of the map. A zero valued α creates no distortion because any number raised to zero equals one. The magnitude of α determines the amount of distortion. An example displacement map after distortion is depicted in Figure 4-3. It should be noted that a pincushion displacement map yields a barrel effect image, and a barrel distortion map yields a pincushioned image. This is the result of the displacement map defining the mapping from the sensor to the scene. In the case of the pincushion displacement map, the inward bowing in scene coordinates of the regularly spaced sensor pixels creates an outward bowing of the resultant image.

A performance gain can be achieved if the geometric distortion is precomputed for the specified UUT and stored in a displacement map. This is used as an initial

81

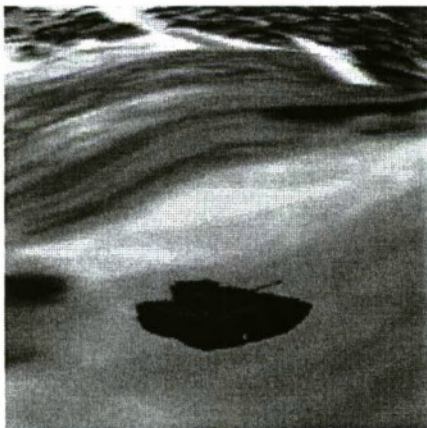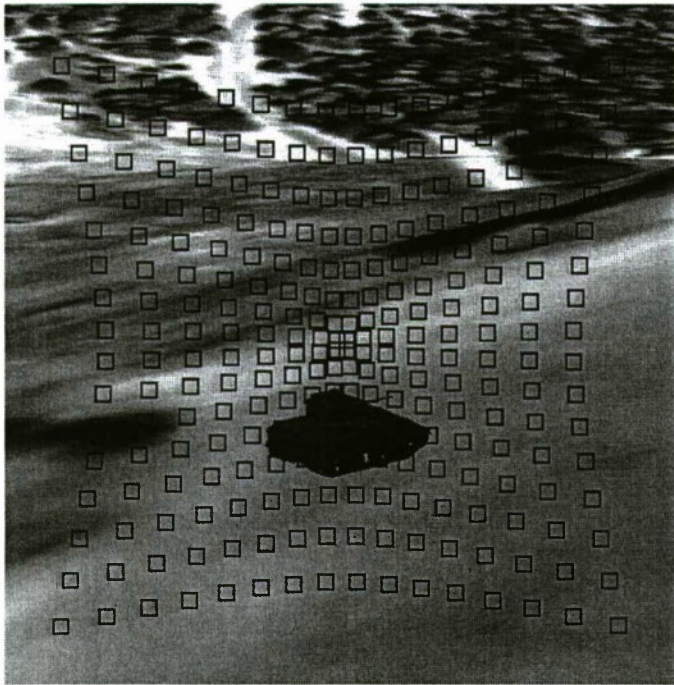displacement map prior to roll and shift latency compensation.

Figure 4-3. Pin-Cushioned Displacement Map

4.3 Other Displacement Functions

Sensors can experience physical jitter when mounted on a moving platform. For a scanning system, this could lead to small perturbations in the location of each successive scan line. If jitter affects the system performance, it should be accurately modeled. Assuming the jitter function can be mathematically modeled, it too can be handled by the UPI. Potential jitter models can include sinusoidal or random per-scan-line or per-pixel effects.

Precomputed jitter offset values are sent to the convolution subsystem and added into the pixel displacement processing.

Displacement processing permits the implementation of any user-defined function that yields new x, y pairs based on the current x, y. Such user-defined functions can be combined with the other displacement effects. However, care should be taken in the order that the functions are executed. For example, functions that are not rotationally invariant must be applied before the rotational latency compensation function is applied. Additionally, the computational complexity of the function should be considered because these operations are performed on a per-sensor pixel basis.

4.4 Displacement Accuracy

The IR/EO system signal processing algorithms can be sensitive to small pixel intensity variations. This technique as described above allows for displacements with scene pixel accuracy; however, this may not be sufficient depending on the application. A solution is to increase the over-sampling factor used in scene generation. This process will increase the subsensor pixel accuracy, but it is limited by the computational resources of both the SGS and UPI. Therefore, a method is presented for displacements with subscene pixel accuracy.

Although convolution is often meant to blur objects, it is typically not intended to displace an object by a fractional scene pixel. However, subscene pixel displacement is desirable for displacement accuracy. Therefore, the goal is to generate a set of kernels that intentionally produce controlled subscene pixel displacements. A kernel is generated by sampling the point spread function (PSF) at intervals equal to the scene pixel width. Since the PSF can be considered a continuous function, it can be resampled starting at any chosen subscene pixel offset, allowing for control of subscene pixel displacements.

Ideally, an infinite number of kernels could be generated to account for the infinite number of subscene pixel offsets, but this obviously is not feasible. Therefore, an analysis was performed to determine how the number of kernels, n_{k}, affects the maximum error due to displacement. Two identical signals were generated: the first, the baseline, was generated with a subscene pixel shift; and second, signal without. A convolution kernel was chosen for the second signal shifted to match the shift between the two signals. Convolution was performed on both signals and the maximum difference between the resultant signals was calculated. This process was

repeated with a varying number of kernels and a plot of the maximum error was generated as shown in Figure 4-4. The subscene pixel resolution is an inverse function of n_k; therefore, the error dropped off inversely with an increase in n_k as expected. Although the decrease in error is limited by an asymptote, it should be noted that this is maximum error calculated over a wide range of subscene pixel offsets. Some offsets produce less error for the given n_k, but none produce more than the plotted results.
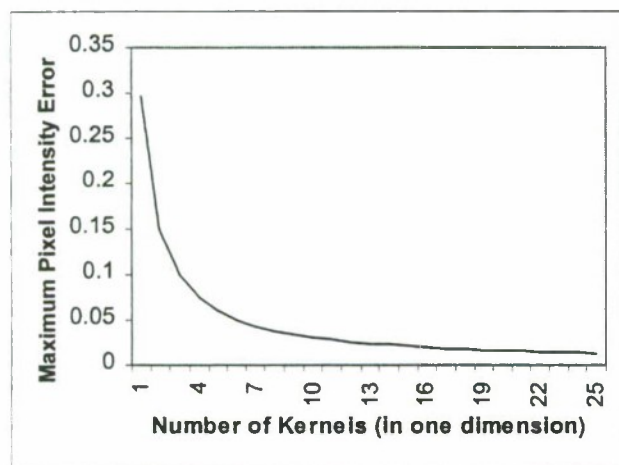
Figure 4-4. Maximum Error as a Function of the Number of Kernels

5. PIXEL EFFECTS MODELING

A variety of noise sources and response nonlinearities that are present in sensors can affect system performance. When performing direct signal injection, the sensor is bypassed. Therefore, these pixel effects are modeled, based on user specified sensor parameters, in the pixel processing subsystem to add to the validity of the HWIL test.

Some of the pixel effects that have been modeled include conversion from radiance to photons, various noise sources, linear responsivity, and automatic gain control (AGC). The convolution subsystem outputs imagery in radiance units. Since most sensors measure incident photons, a conversion factor is applied. This is followed by the introduction of various noise sources, both multiplicative and additive. Linear responsivity models the system's analog-to-digital (A/D) converter by converting the photon count to fixed point units measurable by the A/D. Finally, the AGC normalizes the display to the intensity range of the image.

Since the pixel effects are implemented as software executing on the general-purpose processors in the pixel processor, a wide variety of effects can be modeled in the UPI. User-defined effects, within the limits of the UPI, can easily be added with no change or cost in hardware.

Optical projection can also benefit from the pixel processor. Gain and offset tables can be loaded in the resident memory and used for NUC.

Figure 5-1 displays an example of pixel effects modeling. The first image is that passed in from the output of the convolver. Some pixel effects were applied, resulting in the second image.
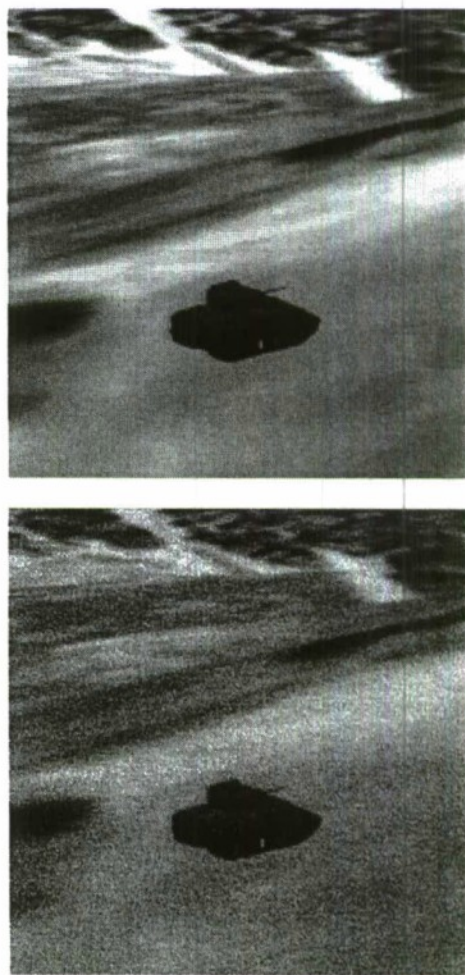
Figure 5-2. Pixel Effects Example6.

6. CONCLUSIONS

The UPI provides a reconfigurable, cost-effective, solution for HWIL IR/EO sensor system testing. A sensor modeling and analysis tool, based on FLIR92, provides the ability to model sensor blurring and per-pixel effects. Sensor parameters can be edited to model these effects in a wide range of current and future sensor systems. Pixel displacement processing can accurately perform latency compensation and can model other sensor effects, including geometric distortion, physical sensor jitter, and various additional user-defined effects. These modeling capabilities are afforded through a combination of both UPI software and hardware. The UPI started as a Small Business Innovative Research (SBIR) project and is being successfully transitioned into a commercial product. It should additionally be mentioned that the UPI has been selected as the Signal Injection Subsystem (SIS) for the joint U.S. Navy and U.S. Air Force Infrared Sensor Simulator (IRSS) program.

7. ACKNOWLEDGMENTS

8. REFERENCES

1. L. Scott, "Analysts Reference Guide," *FLIR92 Thermal Imaging Systems Performance Model*, U.S. Army NVESD report RG5008993
2. P. A. Acevedo, B. E. O'Toole, and T. E. Eisenhauer, "Prototype development of a universal programmable interface for hardware-in-the-loop sensor emulation," *Technologies for Synthetic Environments: Hardware-in-the-Loop Testing III*, Robert Lee Murrer, Editor, Proceedings of SPIE Vol. 3368, pp. 354-365, 1998.
3. D. B. Howe, P. A. Acevedo, and M. E. Nuwer, "Hardware-in-the-Loop Sensor Emulation Using a Universal Programmable Interface," *Technologies for Synthetic Environments: Hardware-in-the-Loop Testing II*, Robert Lee Murrer, Editor, Proceedings of SPIE Vol. 3084, pp. 292-303, 1997.

Integration of a Maritime Modeling Capability into the Infrared Sensor Stimulator (IRSS)

Stephen E. Jacobs
Comptek Amherst Systems, Inc.

Samuel L. Kerr
Naval Surface Warfare Center, Carderock Division

Daryl R. Giles
Naval Air Warfare Center, Weapons Division

ABSTRACT

The Office of the Secretary of Defense (OSD), Central Test and Evaluation Investment Program (CTEIP) is tasked with providing a coordinated process for making joint investments in defense test & evaluation (T&E) to offset the challenges presented by declining investments in test assets and increasing test requirements. Under CTEIP sponsorship, the Navy and Air Force are jointly developing three Joint Installed System Test Facility (JISTF) enhancements that are based on dynamic virtual reality simulation technology. The three enhancements are the Infrared Sensor Stimulator (IRSS), Generic Radar Target Generator (GRTG), and Joint Communications Simulator (JCS). These enhancements will provide each ISTF with the capability to simultaneously test multiple avionics and sensor subsystems installed on an aerospace System Under Test (SUT) (e.g. manned and unmanned aircraft) in a ground test environment. The ISTF enhanced test capabilities will be used to evaluate multi-sensor data fusion/correlation and subsystems' interoperability for Infrared Sensors, RADAR, GPS, and Communications and Data Link subsystems.

The IRSS program[1] was previously briefed at the 1997 and 1998[2] GTM&V Conference. This paper addresses the integration of a maritime modeling capability into the IRSS Scene Generation Subsystem (SGS).

The IRSS system is designed to function primarily on commercial-off-the-shelf (COTS) hardware such as the Silicon Graphics (SGI) Onyx2® InfiniteReality graphics computer. The symmetric multiprocessing capability of the SGI Onyx2 computer gives the IRSS system a multi-channel capability for the simulation and rendering of multi-spectral infrared images at high frame rates.

As part of the Infrared Sensor Stimulator (IRSS) development project, Comptek Amherst Systems Inc. is tasked with incorporating the capability to render infrared simulations of the Maritime Combat Environment (MACE). To achieve the requirements associated with
the modeling and rendering of surface ships, and the dynamic nature of the ocean background, Comptek Amherst has integrated modified portions of the US Navy's maritime model IRENE within the MACE structure. As a result, IRSS has been provided with the unique capability of rendering surface ships and an ocean background in a real-time high-fidelity IR simulations.

This paper will outline the basic process of the integration of the necessary reworked portions of IRENE into IRSS and several of the challenges, issues, and solutions that accompanied this task.

KEYWORDS: Installed Systems Testing, Infrared Sensors, Scene Simulation, Maritime IR Model

1. INTRODUCTION

When creating the IRSS, it became apparent that one of the first sensors to be tested with the system would be the US Navy's AAS-44V FLIR system. In order to ensure that the modeling and simulation capabilities of IRSS would be capable of simulating the type of environment in which the AAS-44V would be used, a requirement was issued for a comprehensive Maritime Combat Environment (MACE) capability to be implemented. There were three main phases to the introduction of a MACE capability into IRSS.

First, the models scheduled for integration into IRSS had to be evaluated for their applicability to MACE. When these models were found to be overall lacking in capability, the second phase became the selection of a suitable model, or set of models, for integration. Third, the selected models had to be integrated into the system.

After briefly reviewing the first two steps outlined above, the paper will focus on the third phase of this process: the actual integration of the selected model(s).

2. IRSS MODELING CAPABILITY EVALUATION

When examining whether or not the capabilities of IRSS meet the challenges of MACE, an examination of some of the seminal features of the maritime environment must be made.

The signature modeling of a surface ship can be adequately described by a one dimensional thermal model. This is due mainly to the fact that most of the outside surface of a ship (hull and superstructure) are composed of relatively thin plates. This fact is made even more relevant by the fact that this plate thickness is very small compared to the exposed surface area of these plates. Therefore, thermally, the plates can be treated as one-dimensional, with almost all of the heat transfer being done perpendicular to the main faces of the plates (internal and external).

The speed of a ship is comparable to most wind speeds encountered. This, coupled with the open-air environment of the ocean, make an accurate atmospheric convection calculation important for a good maritime signature model. Additionally, there are many regions of the ship which have very different temperatures during normal operation. The ability to model these "hot regions" within the structure of a ship is critical to correctly modeling the thermal signature of the ship.

The water environment also presents a very unique and special background for a thermal model. As the background provides a significant amount of radiance to a ship, the background must be modeled as accurately as possible.

The ground vehicle model PRISM and the aircraft model SPIRITS are scheduled for integration into the IRSS. These two models were chosen because of their widespread acceptance in the IR modeling community, and their satisfactory performance in regard to functionality. Both models excel in their respective areas of modeling, however they are not suitable for MACE. PRISM is a detailed 3D thermal signature model used for temperature/radiance calculations of ground-based vehicles. In this modeling realm, a three-dimensional conduction/convection model is necessary due to the thick nature of the plates making up the exterior of the vehicle. While this feature is certainly not something that is bad in a one-dimensional thermal situation, it is unnecessary.

PRISM does not account for the background environment in a way satisfactory for the maritime environment. The method it employs is one of a general background radiance, which is insufficient when dealing with the large effects of solar glint on the sea surface.

SPIRITS is also not up to the maritime modeling task. The model deals with aircraft flying through the atmosphere, not a condition where wind-driven thermal convection is critical. Consequently, the convection modeling capability of SPIRITS is somewhat lacking. Also, the manner in which internal thermal regions are handled is not very transparent, a key issue as the feature is very important to maritime modeling.

Given these considerations, it was decided that investigation of existing maritime signature models had to be undertaken to determine whether an existing model could be incorporated into IRSS, or if a new model would have to be developed from scratch.

3. SELECTION OF A MARITIME MODEL

The investigation of existing maritime models led to two main candidates upon which to base the modeling capability: NTCS/SHIPIR and IRENE. The former was created and maintained by W.R. Davis Engineering, Ltd. Of Ontario, Canada. This model not only calculates the thermal and radiometric properties of a ship in its environment, but it also is capable of simulating complete missile engagements between the ship and a variety of anti-ship missiles. The features important for a complete maritime model are included in this model, making it an excellent candidate. In addition, this model has been accepted as the standard NATO maritime model under RSG.5.

The second model considered, IRENE, is a model developed and maintained by the US Navy at its Naval Surface Warfare Center (NSWC) Carderock Division in Bethesda, Maryland. It also contains the salient features necessary in a signature model for applicability to the maritime environment

Both models are excellent candidates, and either model would have been sufficient to accomplish the MACE modeling goals. The decision was based on integration ease and the fact that the time available for the work to be done was relatively short. While there would not have been a problem working with the team at W.R. Davis, there would have been more issues to resolve with SHIPIR than with IRENE. Also, as IRENE is a US Navy model developed under a US Navy contract, greater control was available to influence the integration effort.

4. INTEGRATION

4.1 INTEGRATION SCHEME

4.1.1 Phase I

The actual integration scheme for MACE was a three-part effort (see figure 1). The first phase involved most of the work on the creation of the new signature module code

itself, and was carried out at NSWC Carderock. A new thermal signature module, derived from IRENE, was written for IRSS. In the past, the model had accepted 2 files for each ship it processed – a geometry file and a separate file containing the radiometric properties of the paints and materials used on the ship. For this effort, these two files were concatenated into one file. This work made the information needed by IRSS more accessible.

The second part of this phase was to make background data (ocean surface radiance values) available to IRSS. This involved the implementation of a way to make the background radiance calculations of the background model consistent with the data used to thermally process the ship.

A major concern was the method to be used in rendering the ocean background. One feature of most maritime models is that while the ocean surface would be accounted for in the radiometric calculations, the information would not be retained for rendering. This background information is needed by IRSS.

A separate ocean surface background model based on the model KELSEA, developed by the Naval Research Laboratory, is used for rendering backgrounds. This model uses a ray-tracing algorithm to determine the apparent radiance of the sea surface based on look angle, sun position and observer altitude. This model was modified to write out the sea surface radiance texture maps (SSRTMs) which are used by IRSS.

The actual rendering of the background is done through the use of the IRSS radiance map rendering capability. The radiance of the ocean surface is calculated in a square grid fashion. The squares comprising this grid can be either 1m or 5m on edge and are 512 x 512 grids. The resultant radiance maps are therefore 512m x 512m or 2560m x 2560m in size, respectively. These grids are then tiled across a flat-polygon ocean surface background appropriately, depending on the size of the background used.
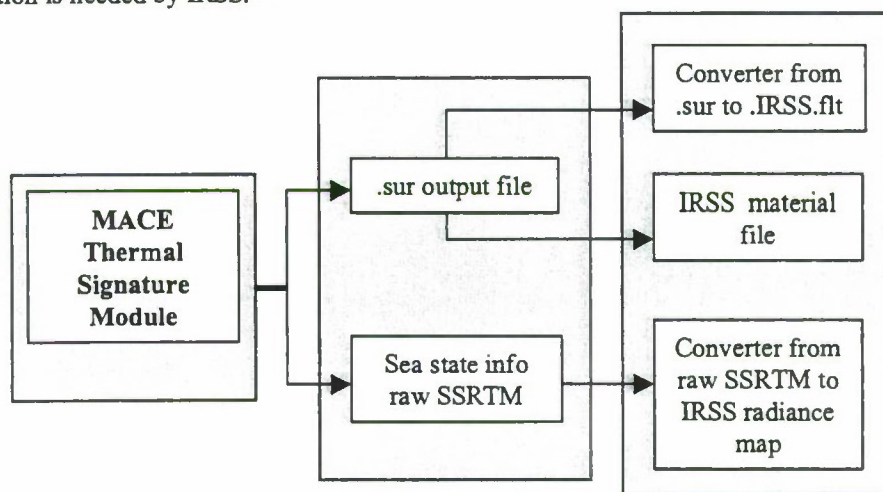
Figure 1. Basic IRSS/MACE Integration Scheme.

4.1.2 Phase II

As shown in figure 1, the next part of the overall effort (carried out at Comptek Amherst Systems in Buffalo, NY) involved the creation of a file converter. The converter extracted the necessary information from the .sur files (the modified IRENE output file) and entered them into .IRSS.flt (OpenFlight) format, which is the three-dimensional geometry format used in IRSS. This work was facilitated by the use of the Application Programmable Interface (API) provided with Multigen-Paradigm's modeling program. The API allows for fast

and easy access to the OpenFlight file format, making the conversion a relatively simple task.

Also part of Phase II was the conversion of the radiance data produced by the KELSEA-based sea surface model into a format usable by IRSS. The radiance data was in a 16-bit, 256 color grayscale format, which was converted into Silicon Graphics .bw texture format. This texture was then rendered as a radiance map, with atmospheric attenuation and path radiance calculated on a pixel-by-pixel basis.

A related task was the creation of a material file for IRSS, containing the radiometric properties of the paints and

materials of the ship to be rendered. As previously stated, IRENE at one point used a separate file for the materials used on the ship. This was incorporated in the actual geometry file of the ship itself. This file is processed by the IRSS converter, which in turn creates a geometry file in OpenFlight format and a material file for each material and/or paint used on the ship.

4.1.3 Phase III

Phase III involved some of the first work in the project, but it was actually finished last. It consisted of the complete rewriting of the IRENE thermal module, making it, in effect, a completely new thermal model for inclusion in IRSS.

4.1.4 Phase IV

Not depicted in figure 1 is Phase IV of the integration effort. This part mainly involved the creation of a user-friendly front-end for MACE. The current method of user interaction with the MACE thermal signature module is a call through a UNIX command-line script. This script processes the desired atmospherics and environmental information, selects the ship model to be processed, and calls the actual signature module. While workable, a more user friendly graphical interface has been proposed for the fourth integration stage. Currently, it is considered a future upgrade issue. It is important to note, however, that there is no functionality lost by not completing this work, as it is basically an ease-of-use issue.

4.2 EXECUTION OF INTEGRATION SCHEME

The work at NSWC Carderock (Phases I and III) went very well. The Carderock team stayed in close contact with the Comptek Amherst team, making it easy for any issues that arose during the work to be resolved quickly and to everyone's satisfaction. This close cooperation led to the saving of a great deal of time and effort, as little time was spent doing work that was later thrown away. This made the reworking of the input file format and creation of the new signature module very smooth.

Phase II was similarly straightforward for the Comptek Amherst Systems team. The task of creating the file converter, as mentioned above, was completed very quickly due to the fact that the MultiGen's *Creator* Plug-In Tools API gives the programmer fast and easy access to the OpenFlight format.

5. RESULTS

Depicted in figures 2 through 7 are some examples of images rendered by IRSS using MACE-generated ships and ocean surface textures. The exact details of each screen capture's contents are given in the caption for each For all images, the sensor is a 256 x 256 staring LWIR array with a 10° x 10° field of view, and white is hot.

In general, the results are very satisfactory. The ships displayed in the images are actually two renderings of the same ship facing in nearly opposite directions. This was done to illustrate (within the limitations of graphical reproduction) the effect of sun position on thermal signature of the ship. The atmosphere used in the rendering is a clear atmosphere – therefore there is little attenuation of the targets and ocean surface texture as a function of distance from the sensor.

A few important notes should be made about the ocean surface texture. Each texture is created based on a single look angle. Naturally, during the course of a scenario, the look angle of the sensor can change dramatically. This limitation is an important one, as it places a limit on the realism of the ocean background. A dynamically updated series of textures is a possibility for future capability, as this would increase the accuracy of the rendered image. This need can be through appropriate extensions to the texture paging scheme of IRSS.

A fact worth mentioning is the almost uniform appearance of the ships in these figures, aside from the solar heating. There was no attempt to create accurate hot part information for these particular ships as sensitive data would be needed for such a task. Tests indicate expected results when using realistic heated compartments in a ship model.
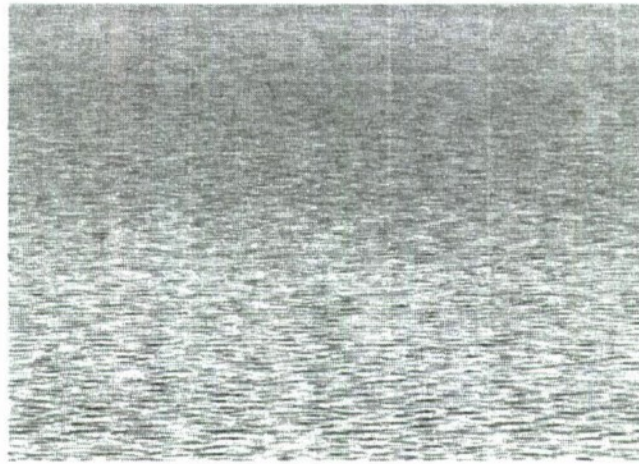
Figure 2. This image is a snapshot of the ocean surface texture at 1m resolution and a sea state of 2. The sensor altitude is ~1200m with a look angle of –10°.

Figure 3. The ocean surface texture is 1m resolution and sea state 2. The sensor altitude is ~1225m with a look angle of –5°. Ranges to the two ships are 2650m and 3600m.

Figure 4. The ocean surface texture is 1m resolution and sea state 2. The sensor altitude is ~1365m with a look angle of –5°. Ranges to the two ships are 1200m and 2150m.

Figure 5. The ocean surface texture is 1m resolution and sea state 2. The sensor altitude is ~1400m with a look angle of −10°. Ranges to the two ships are 740m and 1560m.

Figure 6. The ocean surface texture is 5m resolution and sea state 2. The sensor altitude is ~1225m with a look angle of −5°. Ranges to the two ships are 2650m and 3600m.
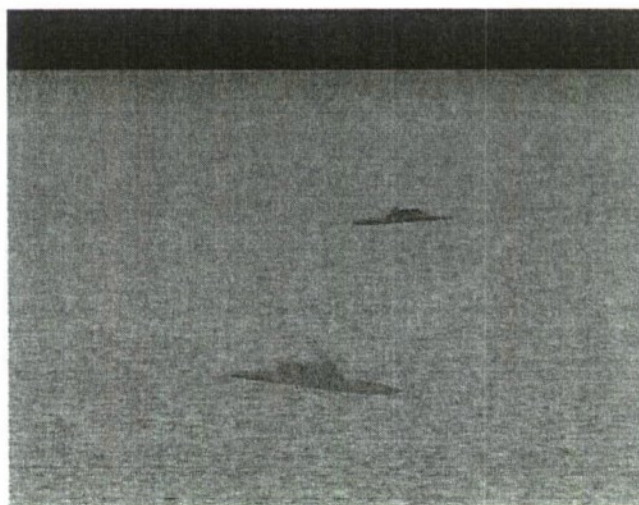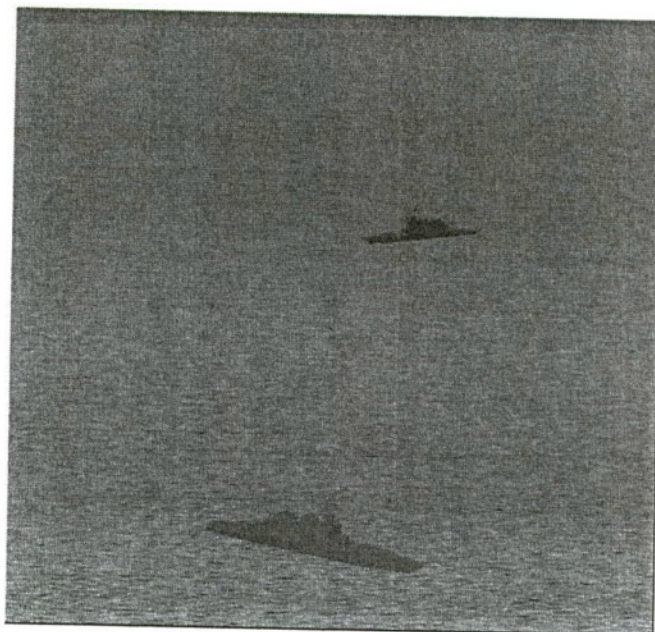
Figure 7. The ocean surface texture is 5m resolution and sea state 2. The sensor altitude is ~1365m with a look angle of −5°. Ranges to the two ships are 1200m and 2150m.

6. SUMMARY AND FUTURE

Many of the important issues involved in the integration of a Maritime Modeling Capability into IRSS have been reviewed. There are several lessons which can be drawn from this success, as well as some important points for further work and improvement.

The use of MultiGen Paradigm's *Creator* for 3D modeling by IRSS allowed the conversion of the .sur geometry format to be extremely smooth. This was critical due to the fact that the entire project, from initial funding to final implementation, lasted only 5 months.

Another important point relates to the cooperation between the developers of IRENE and Comptek Amherst Systems. They were available and willing to participate in the effort. The work would have been far more difficult, and certainly not nearly as straightforward, without their participation. With their knowledge of the workings of an already validated maritime model, the work to be done in the creation of the new signature module was completed very quickly.

There is still room for improvement in the MACE modeling capability of IRSS and studies in this direction have already been undertaken. Features earmarked for future effort are the creation of a wave-height algorithm for sea surface rendering and the inclusion of plumes on the ships themselves. These features have been identified as the main enhancements required to bring the MACE modeling effort closer to its desired level of functionality.

7. REFERENCES

[1] O'Toole, B. E., "Real-time Infrared Scene Simulator (RISS)," Robert Lee Murrer, Jr., Ed., *Proceedings of SPIE*, **2741**, 209-218, Orlando, FL, 1996.

[2] Joyner, T., "Real-time IR/UV EO Sensor Stimulation," *Proceedings of the Ninth Annual Ground Target Modeling & Validation Conference*, 166-180, Houghton, MI, 1998.

DEGRADING IMAGES FOR REALISM

R. M. Geatches and G. J. Bishop,
Sowerby Research Centre, British Aerospace,
Filton, Bristol, BS34 7QW, England.

ABSTRACT

The Sowerby Research Centre (SRC) based at Filton, Bristol is responsible for developing infra-red signature prediction capabilities for it's parent body, British Aerospace. To achieve this, the *SIRUS* code has been developed and used on a variety of projects for the last decade. *SIRUS* is capable of providing infra-red predictions on air breathing and rocket motor propelled vehicles, with considerable accuracy.

To enable *SIRUS* to remain both flexible and efficient for the ever increasing prediction requirements of the various business units within BAe, a suite of pre- and post-processing codes have been developed. With a view to performing true stealth design and assessment, SRC have recently developed the *DEGRADE* post-processing codes.

DEGRADE allows for the *SIRUS* output images to be post-processed in such a way as to produce more realistic imagery. The *SIRUS* images are composed of a target embedded in a homogeneous background. *DEGRADE* can take this target image, and embed it into an alternative background. This background can be measured or simulated, and *DEGRADE* can produce both these requirements, by conversion of measured data to a compatible format or by simulating backgrounds using Gaussian Markov statistics. Furthermore, *DEGRADE* can apply an MTF and noise, to simulate how the image would be seen as viewed through a sensor. Alternatively, the image can be transformed to a format compatible with more sophisticated sensor modelling codes, such as the TACOM Thermal Image Model.

This paper will describe how *DEGRADE* can be used to take simulated *SIRUS* images and create more realistic images, that can be used for better stealth design and assessment.

1. INTRODUCTION

Low-observable vehicle design is a key factor in retaining the element of surprise during conflict and for improving mission survivability. These considerations are not limited to the radio-frequency, which traditionally is perceived to be the main stealth consideration. Increasingly, the total signature from all frequencies across the electromagnetic spectrum must be considered in the overall design of a low-observable vehicle. Predicting Infra-red (IR) signatures has therefore become an important design feature, and over the last decade BAe has developed the *SIRUS* code to achieve this.

SIRUS is capable of providing IR signature predictions on air-breathing and rocket propelled vehicles with considerable accuracy. The properties that are modelled by *SIRUS* include:

- Surface temperature

- Surface reflectance (parameterised Bidirectional Reflectance Distribution Function)

- Cavity physics

- Plume gas radiative transfer (including particulate scattering)

- Atmospheric effects (including solar contributions)

- Backgrounds

- Sensor effects (imaging and threshold detection)

A key element in the design process however, is not necessarily the IR signature itself, but the contrast with its environment. It is quite feasible that the IR signature of a vehicle may have components that are in positive contrast to a particular background, and others that are in negative contrast. This could render the IR signature contrast-neutral when unresolved in the imager (sub-pixel). Whereas, when the aircraft is closer, thereby imaged by multi-pixels, the IR signature structure is resolvable and the neutral contrast seen at sub-pixel resolution is lost. Furthermore, the structure of the IR signature, in some circumstances may improve the low-observability at resolved distances, as the structure may act as a camouflage against a cluttered background.

To enable *SIRUS* to remain both flexible and efficient, a suite of pre- and post-processing codes have been developed alongside. One such code, *DEGRADE* allows

predicted signatures to be embedded into a variety of alternative backgrounds, against which their observability can be assessed. The *DEGRADE* suite of code can also create and/or convert these backgrounds, and transform the images to simulate how the image would be seen as viewed through an IR sensor. These applications add realism to the signature predications provided by *SIRUS*, and this paper describes how *DEGRADE* can be used to enhance low-observable design and counter-stealth assessment.

2. BACKGROUNDS

There is a necessity within the Aerospace industry to create representative models of backgrounds for use with our *SIRUS* target IR prediction capabilities. There are software packages available to generate synthetic backgrounds, however, these tend to rely intensively on skilled users. Furthermore, it is not possible to generate nor measure every conceivable background scenario required. Instead, BAe have adopted the approach of generating most of the backgrounds required using Gaussian Markov statistics. The *DEGRADE* suite of codes has the capability of generating such backgrounds, as well as converting measured imagery to a format compatible to be used with the *SIRUS* target images.

The degree of complexity used in generating backgrounds is, to a certain degree, a matter of user choice. The simplest method is to use a homogeneous background, where a constant value (*e.g.*, radiant intensity, $I[WSr^{-1}]$) is used for every pixel in an image designated non-target. The next level of complexity would be to allow a statistical variance of the mean value, so that the radiant intensities of the background pixels vary across the image, *e.g.*, $\bar{I} \pm \sigma[WSr^{-1}]$. The method employed here is one level of complexity greater, whereby the background pixel values are generated using a statistical variance of the mean plus a correlation length. To do this, Gaussian Markov statistics have been employed, and forms part of the *DEGRADE* suite of *SIRUS* post-processing codes. A variety of statistical parameters across an image may also be generated by manually combining separately generated files.

In this way, it is possible to use *DEGRADE* to generate a wide range of statistically representative backgrounds for use with *SIRUS* target images. It then becomes very simple to generate multiple background images that can be used with the *SIRUS* images, to enable the low-observability to be assessed for varying environments. Deciding on which statistical parameters to use to generate these backgrounds is more difficult. A

good starting place is often to statistically assess a suitable measured background. Given a file containing pixels of measured apparent temperature (as for AGEMA 900 measurement files) the mean ($\bar{T}$) and variance (σ^2) can easily be determined by:

$$\bar{T} = \frac{1}{N} \sum_{n=1}^{N} T(n) \tag{1}$$

$$\sigma^2 = \frac{1}{N-1} \sum_{n=1}^{N} (T(n) - \bar{T})^2 \tag{2}$$

where, $T(n)$ are the apparent temperature values of the pixels and N is the total number of pixels. The calculation of the variance, σ^2, sometimes uses $1/N$ rather than $1/(N-1)$. The small sample approximation using $1/(N-1)$ is used in *DEGRADE*.

To obtain a representative correlation length, the image is assessed by passing a linear mask over the pixels, in order to calculate the mean auto-covariance function:

$$\overline{C(l)} = \frac{1}{N} \sum_{n=1}^{N} (T(n) - \bar{T}) \times (T(n+l) - \bar{T}) \tag{3}$$

where, l is the separation distance between pixels (see figure 1, hence, for a horizontal mask of 2 pixels $l = 1\times$ pixel width, and a vertical mask of 3 pixels $l = 2\times$ pixel heights).
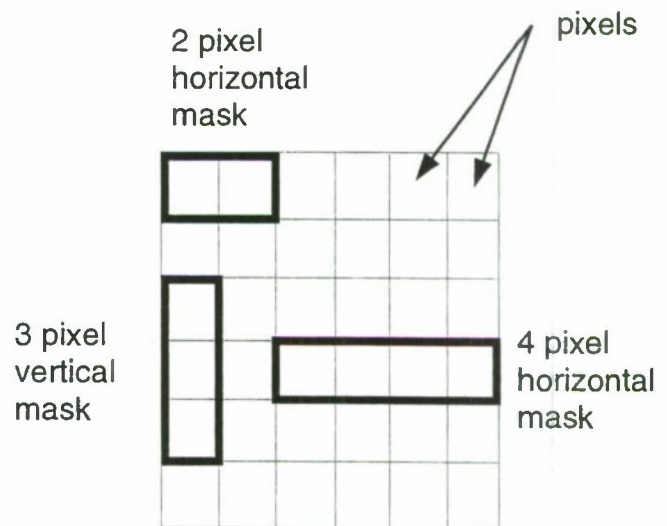

Figure 1: Examples of the horizontal (x-) and vertical (y-) masks used to assess the statistics of a pixelated image.

The correlation length, $\alpha^{-1}[m]$ can then be calculated from:

$$\overline{C(l)} = \sigma^2 e^{-\alpha l} \rightarrow \alpha^{-1} = \frac{-l}{ln(\overline{C(l)}/\sigma^2)} \qquad (4)$$

An example of an image measured by an AGEMA 900 is given in figure 2.
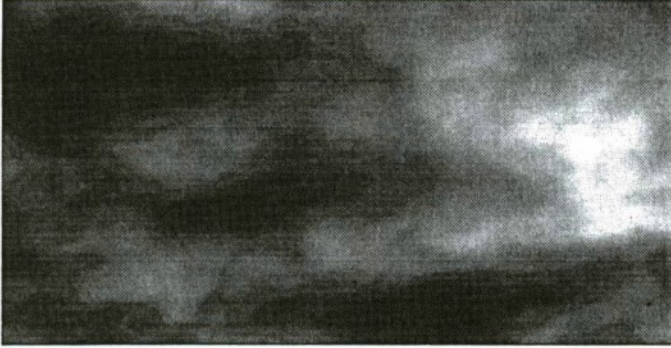

Figure 2: Measured scene using an AGEMA 900. Scale 1.293-2.579 W/Sr.

The image is a sky scene showing some cloud structure. The statistics of this image were assessed, as described above. These values were then used to generate an image using the Gaussian Markov statistical method, and the resulting image similarly assessed. Table 1 gives the results of these statistical analyses, and figure 3 illustrates the resulting statistically generated image.

Figures 4 and 5 indicate the distribution of temperatures and auto-covariance functions for the measured and generated images, respectively. These figures, together with the results shown in Table 1, indicate that the images compare very well. Furthermore, the images themselves (figures 2 and 3) show distinct similarities.

TABLE 1: Statistical assessment results of the measured and simulated images.

	Measured image	Simulated image
$\overline{T}[K]$	264.52	264.52
$\sigma^2[K]$	41.53	41.06
$\overline{\alpha}^{-1}[m]$	52.16 ± 8.60	46.86 ± 7.68

The statistical assessments and image simulation are achieved using apparent temperature values $[K]$. Radiant intensities $[W Sr^{-1}]$ are more appropriate values to use, as they are independent of pixel area. Calculation of
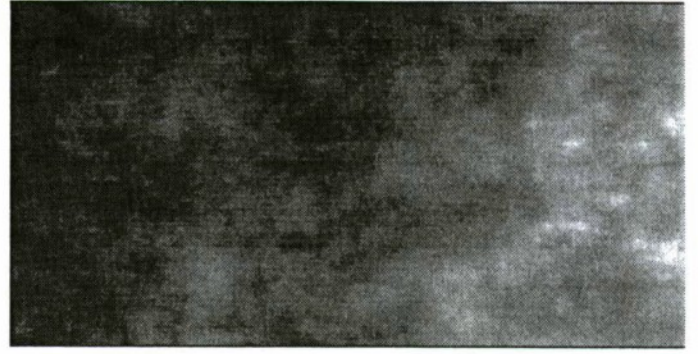
Figure 3: Simulated image using the statistical assessment of the measured image. Scale 1.293-2.579 W/Sr.

radiant intensity from apparent temperature is achieved by Planck function interpolation. However, this assumes that the simulation range between sensor and target is constant for all pixels. If this is not the case, then errors will result. A pixel range map would allow this potential problem to be avoided, and this is planned in future developments of *DEGRADE*.

3. EMBEDDING A TARGET

One of the main tasks of *DEGRADE* is to take a *SIRUS* image file, extract the target from the homogeneous background and embed the target into an alternative background. It is assumed that the *SIRUS* target will be a high resolution image with the background at a lower resolution. The embedding process is achieved by area-weighting the radiant intensity values of the target and background pixels, thereby creating a new image. The area-weighting of the radiant intensities is achieved thus:

$$I_{res} = I_b \frac{(A_b - A_{t'})}{A_b} + I_t \frac{A_{t'}}{A_t}$$

where, I_b is the radiant intensity, and A_b is the area (cm^2) of the background (this is equivalent to the projected pixel area); $A_{t'}$ is the target area component residing within the background; and I_t is the radiant intensity, and A_t is the area of the target pixels. This calculation is performed for each target and background pixel, hence calculating a resulting radiant intensity, I_{res}.

94

Figure 4: Temperature distribution of the measured and simulated images.

Figure 5: x- and y- auto covariance functions of the measured and simulated images.

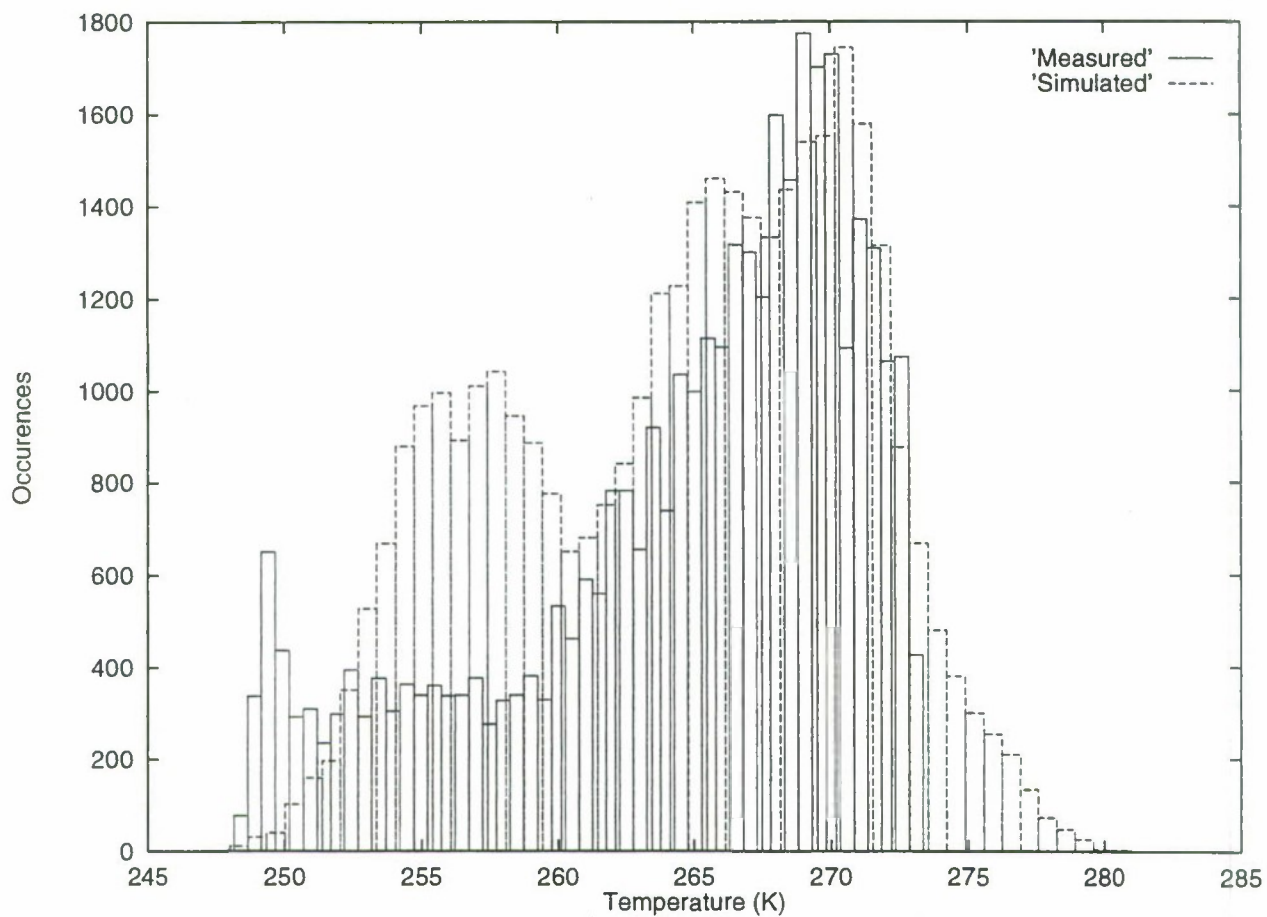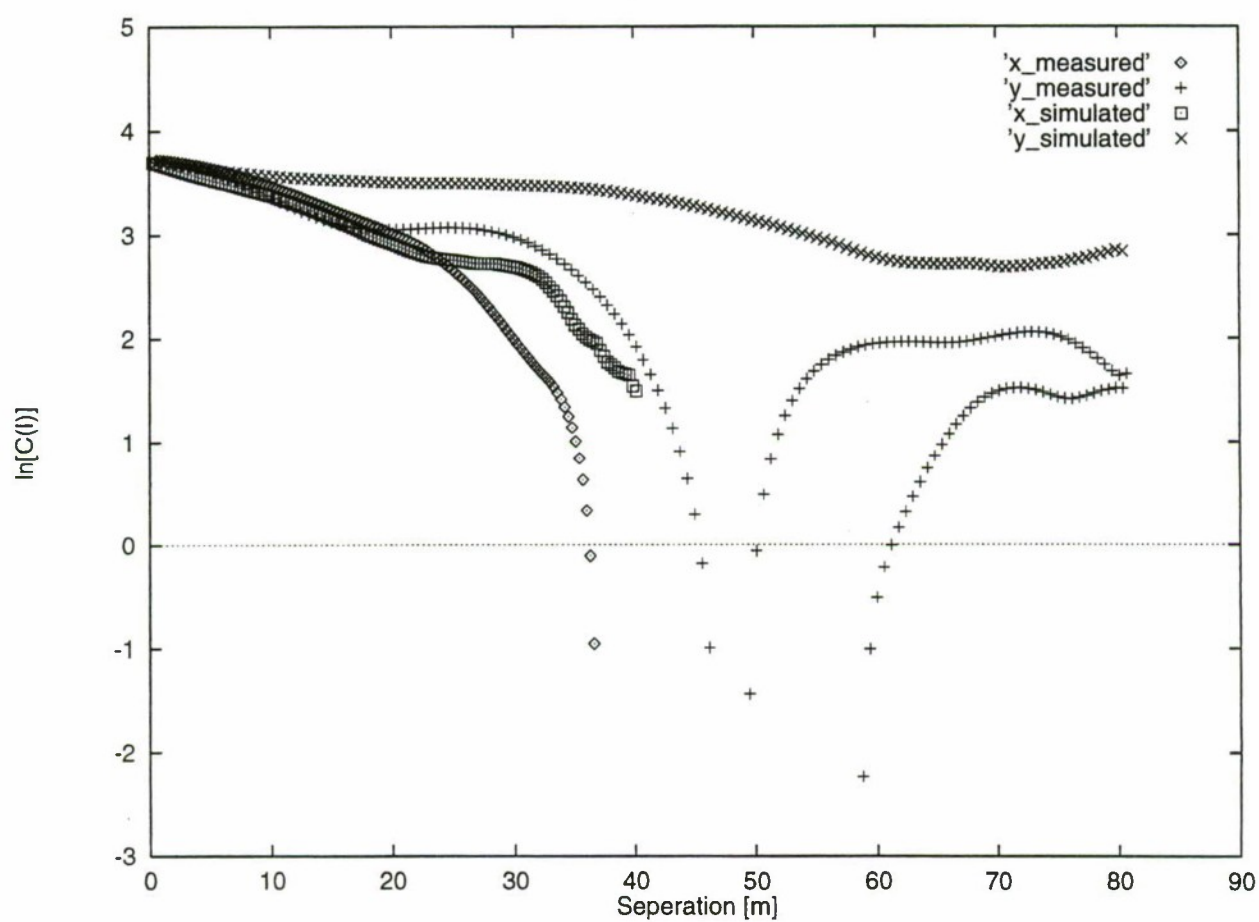The image in figure 6 indicates the resulting image of a target embedded into the *DEGRADE* simulated background, shown in figure 3. The image in figure 7 shows the same target, in the same position embedded in a homogeneous background.

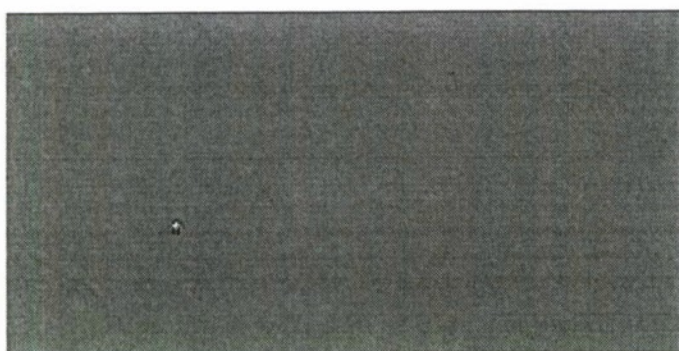Figure 6: Target image embedded into a cluttered background.

Figure 7: Target image embedded into a homogeneous background.

Figure 8: Target image embedded into a cluttered background, with sensor effects (line spread function and white noise) added.

The radiant intensity of the homogeneous background was taken as the mean radiant intensity of the target, so that the image would be contrast neutral if unresolved (sub-pixel). The target is clearly resolvable, spread across multiple pixels (37 target pixels in a total of 36450). From figure 6 it can be seen that the target in the cluttered background is well camouflaged by its surroundings. Whereas, the target embedded into a homogeneous background (figure 6) is clearly distinguishable from its surroundings.

4.   ADDING SENSOR EFFECTS

The images shown in figures 6 and 7 are indicative of the effect of clutter on the ability to distinguish a target from background. However, these images are idealised cases, as no account of the sensor response has been included. *DEGRADE* can take these images and apply simple sensor effects, in terms of a line spread function and white noise. Figure 8 shows the image given in figure 6, with these simple sensor effects included.

By adding a line spread function and adding some white noise, the target in the cluttered background becomes even more difficult to distinguish from its surroundings.

To enable the full extent of the sensor effects to be simulated, the *SIRUS* output can be integrated with other codes, such as the commercially available code TTIM (OptiMetrics, Inc., 3115 Professional Drive, Ann Arbor, Michigan, USA). Figure 9 illustrates the capabilities of the *DEGRADE* suite and its potential integration with other codes and facilities.

Image files (from, *e.g.*, *SIRUS*; AGEMA measurements; Statistically generated backgrounds; ACSII data) can be converted into alternative formats for assessment. The *IMAGE* sub-program of *DEGRADE* allows these image files to be converted to a variety of output formats, as shown in figure 9. *VIZ* is a visualisation code, developed by BAe, whereas TTIM is the TACOM Thermal Imager Model, used to evaluate sensor effects. The effects that can be modelled using TTIM are far more extensive than the *DEGRADE* sensor modelling sub-program (*MTF*), and include:

- Atmospheric and battlefield effects
- Detector responsivity
- Atmospheric and optics MTFs
- Optional spatial sampling
- Detector noise
- Post-detection MTFs

97
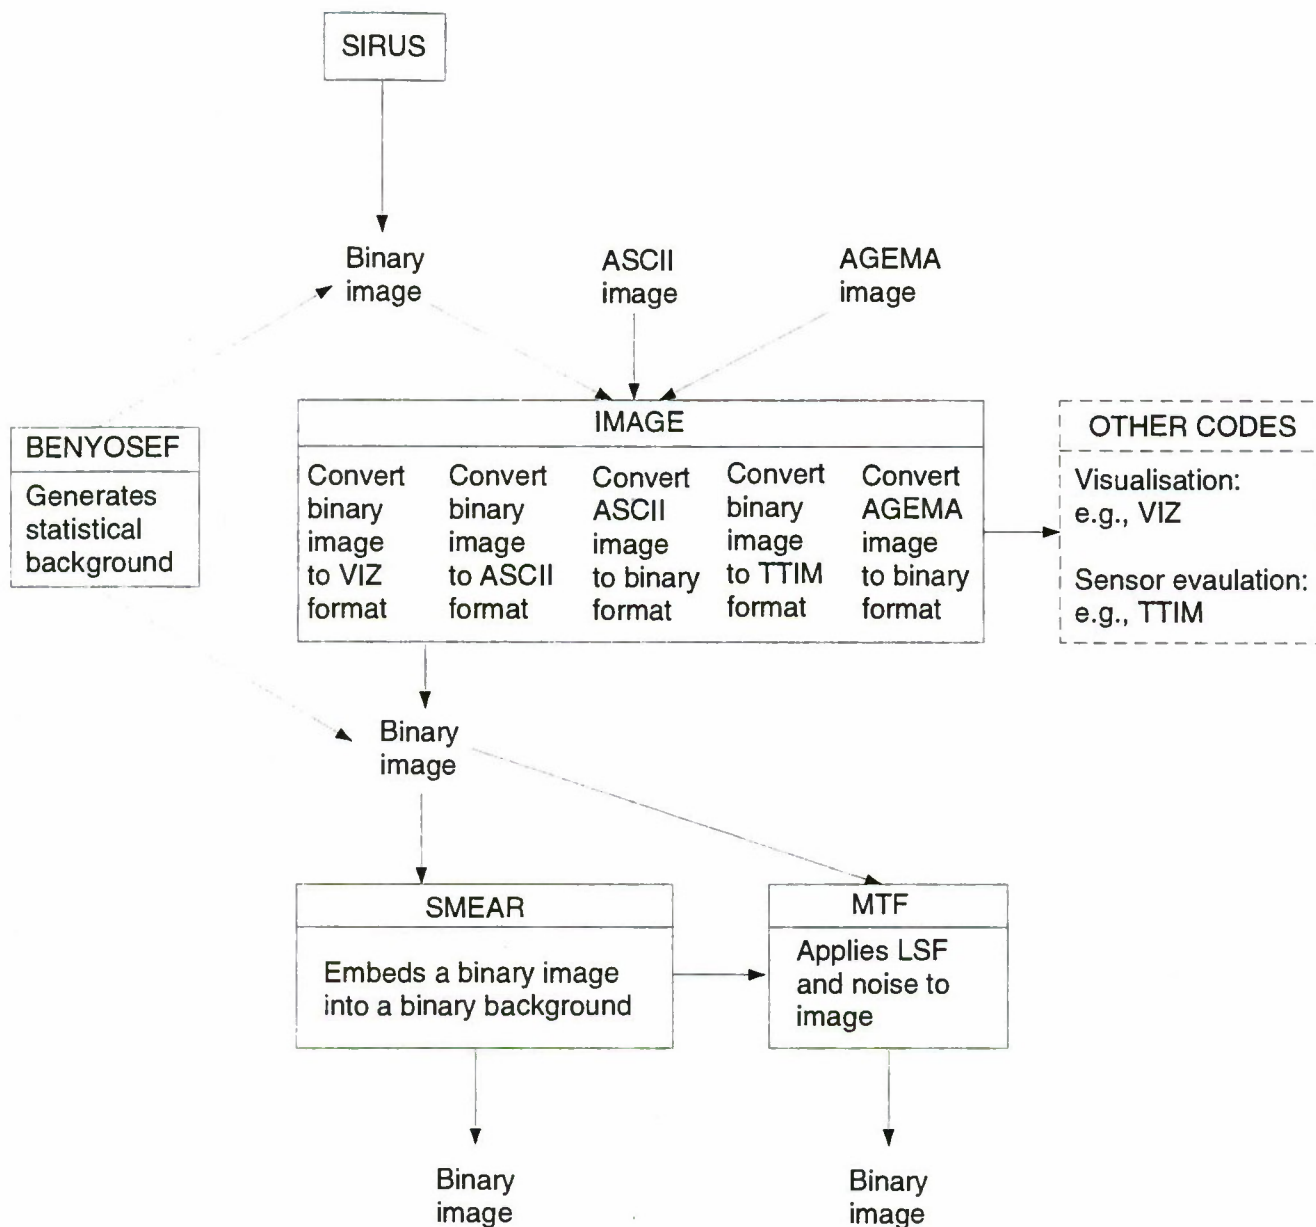
Figure 9: Flow chart illustrating the processing capabilities of DEGRADE.

5. CONCLUSIONS

The British Aerospace code, *SIRUS* in conjunction with its post-processing suite of codes, *DEGRADE* can be used as powerful tools in the design of low-observables. A wide variety of problems can be addressed, ranging from simple point signature assessments through to contrast analyses and probabilities of detection. This end-to-end capability is of use to both the low-observable and the counter stealth designers. *SIRUS* is capable of providing high resolution Infra-red predictions on air-breathing and rocket motor propelled vehicles, with considerable accuracy. *DEGRADE* on the other hand, can take these *SIRUS* images and create realistic scenarios. This is done by generating suitable cluttered backgrounds, embedding the target images into them, and applying sensor effects to simulate how the images would be seen as viewed through a sensor. This then allows for the detectability/effective camouflage of a vehicle to be assessed in a variety of environments.

These capabilities, and the continued development of both *SIRUS* and *DEGRADE*, create a suite of powerful codes that can fully address the requirements of stealth and counter-stealth design.

ACKNOWLEDGMENTS

Exact calculation of reflected earth/sky radiance from a plate with specified BDRF, and with arbitrary orientations of plate and detector

M. J. Caola & N. C. Roberts, British Aerospace Sowerby Research Centre, Filton, Bristol BS34 7QW, U.K.

Abstract The accurate calculation of the I.R. signature of a multi-facet/plate model of a real vehicle in a real environment must include the the correct physical modelling of accurate BDRF reflectance and atmospheric radiance. The measured BDRF of a plate is fitted to the OPTASM model with cubic-spline interpolation, for use in the 2-D integral of the radiance reflected from the plate. The earth/skyshine on the plate is similarly inter-polated from MODTRAN at given height and atmosphere. The radiance reflected from the arbitrarily orientated plate to an arbitrary detector di-rection is calculated numerically using the basic radiance integral with the (now) continuous BDRF and MODTRAN functions. We give numerical ex-amples of this detector radiance in various BDRF-MODTRAN-orientation scenarios. The reflectance is calculated directly in spherical-polar plate co-ordinates, but the results are expressed in the earth co-ordinates of plate and detector .

Analysis

Recent requirements and advances in I.R. stealth technology necessitate ac-curate calculation of the atmospheric radiance reflected from an vehicle sur-face. The surface is usually modelled as a connection of flat facet/plates, so this calculation is the sum of radiances reflected from single plates: each has different area A_m, orientation and surface reflectance chararcteristics. This reflectance is completely described by the surface BDRF (bi-directional reflectance function) ρ', which experimentally is the fraction of radiance re-flected into an arbitrary direction from a (different) arbitrary incident di-rection. The earth's atmospheric radiance is accurately modelled by a code such as MODTRAN, and depends only on the elevation angle θ to the earth surface normal z direction.

The plate m is orientated at elevation θ_m and azimuth ϕ_m to the natu-ral earth axes xyz, and at height h above the earth (we envisage radiance summation over plates $m = 1, 2, ...M$ comprising the model). The plate axes XYZ are fixed on the plate, with Z as the plate normal. Starting with co-incident earth- and plate-axes, the latter are rotated by ϕ_m about z followed by θ_m about Y', which specifies the XYZ axes (see fig.1).

The radiance reflected to an arbitrarily orientated detector is the sum over the 2π sr above the plate surface of the product of atmospheric radiance

with BDRF. This calculation can be done in earth xyz or plate XYZ, but must of course be expressed in 'real-world' earth co-ordinates.

We present here the calculation in plate XYZ, then referred back to earth xyz. We have also done the same completely in earth xyz: exact agreement between the two, involving cases with Helmholtz reciprocity, leads us to believe that both analytical methods are correct.

(Notation: we use upper-case Roman and Greek symbols for plate variables, and lower case for earth variables. Thus
plate: radiance I, I', polar $\Theta, \Phi, \Theta', \Phi'$, direction $\Omega(\Theta, \Phi)$
earth: radiance i, i' , polar $\theta, \phi, \theta', \phi'$, direction $\omega(\theta, \phi)$, plate θ_m, ϕ_m
Here f is an *incident* and g' is a *reflected* variable.)

The plate-system radiance reflected to a detector in spherical-polar direction (Θ', Φ') is [1]

$$I'(\Theta', \Phi', \theta_m, \phi_m) = \int_0^{2\pi} d\Phi \int_0^{\pi/2} d\Theta \ I(\Theta, \Phi) \rho'(\Theta, \Phi, \Theta', \Phi') \cos(\Theta) \sin(\Theta)$$

(1)

The BRDF ρ' is the fraction of incident radiance $I(\Theta, \Phi)$ reflected in detector direction (Θ', Φ'), and is physically modelled by the OPTASM expression [2]

$$\rho'(\Theta, \Phi, \Theta', \Phi') = \frac{A(\Theta)}{B(\Theta) + 1 - \Omega_{pk}(\Theta_{pk}, \Phi_{pk}).\Omega'(\Theta', \Phi')}$$

(2)

where $\Omega_{pk}.\Omega'$ is the dot-product of the two unit-vector directions, *i.e.* the cosine of the angle between them, and Ω_{pk} are the OPTASM-fitted peak direction(s). We consider only an 'isotropic' non-figured plate surface here, so

$$\rho' = \rho'(\Theta, \Theta', \Phi - \Phi')$$

(3)

and

$$\Phi_{pk} = \Phi - \pi.$$

(4)

We also note that (Helmholtz) reciprocity ('reversibility of light') means

$$\rho'(\Theta, \Phi, \Theta', \Phi') = \rho'(\Theta', \Phi', \Theta, \Phi)$$

(5)

Our final reflected-radiance 2-D integral is, referred back to earth detector (θ', ϕ'),

$$i'(\theta', \theta_m, \phi' - \phi_m) =$$

$$\int_0^{2\pi} d\Phi \int_0^{\pi/2} d\Theta \frac{I(\Theta, \Phi) \cos(\Theta) \sin(\Theta) A(\Theta)}{B(\Theta) + 1 - \cos(\Theta_{pk}(\Theta) \cos(\Theta') + \sin(\Theta_{pk}(\Theta) \sin(\Theta') \cos(\Phi - \Phi')}$$

(6)

To practically calculate this integral we need I, Θ', Φ' expressed in earth co-ordinates:

$$I(\Theta, \Phi) = i(\theta) \qquad\qquad \theta = \theta(\Theta, \Phi, \theta_m)$$

101

$$\Theta' = \Theta'(\theta_m, \phi_m, \theta', \phi')$$
$$\Phi' = \Phi'(\theta_m, \phi_m, \theta', \phi') \tag{7}$$

These expressions and the integral are calculated using Mathcad intrinsic functions. Firstly, instrinsic *genfit* was used for the best Levenberg-Marquardt fit of OPTASM $\{A, B, \Theta_{peak}\}$ for each of the discrete in-line BDRF data at incident $\Theta_{inc} = 0, 10, 20..85^o$. Secondly, intrinsic cubic-spline *cspline* and interpolation *interp* were used to create continuous OPTASM $\{A(\Theta), B(\Theta), \Theta_{peak}(\Theta)\}$ and continuous MODTRAN $i(\theta)$ from the above discrete ones. Finally, the continuous $\{i, A, B, \Theta_{peak}\}$ and expressions (7) were used in Mathcad (6).

The detector also receives self-emitted radiance from the plate

$$I'_{self}(\Theta', \Phi') = B(\lambda, T_m)[1 - \rho(\Theta', \Phi')] \tag{8}$$

where B is the Planck function at wavelength λ and plate temperature T_m, and HDR

$$\rho(\Theta', \Phi') = \int_0^{\pi/2} d\Theta \int_0^{2\pi} d\Phi \rho'(\Theta, \Phi, \Theta', \Phi'), \tag{9}$$

which integral is calculated in Mathcad as above.

Results are shown for a BDRF having no sensible specular component, which was thus fitted with a single OPTASM term, as in (2). (A good fit to semi-specular BDRF was also obtained with a two-term model.) We append Figs.2–10 of MODTRAN radiance $i(\theta)$, a sample BDRF fit, and detector radiance $i'(\theta', \theta_m, 0)$ as a function of detector elevation angle θ' for plate elevations $\theta_m = 0, 30, 60..180^o$. We believe that these $i'(\theta')$ are exact reflected radiance under the specified conditions. They generally peak markedly at θ'_{max} and θ'_{min}, due to the rapid increase of both ρ' towards grazing incidence and of skyshine $i(\theta)$ above $\theta \simeq 90^o$. However, the more fundamental reflected power w' at the detector

$$w' = i' \cos \Theta' \Delta \omega_{det} A_m \tag{10}$$

shows less variation.

References

1. Siegel R. and Howell J. R., 'Thermal Radiation Heat Transfer', 3e, Hemisphere (1992)

2. Acquista C. and Rosenwald R., "Multiple Reflections in Synthetic Scenes", Proc. 5^{th} Ann. Ground Target Modelling & Validation Conf.", Houghton (MI), (1994)
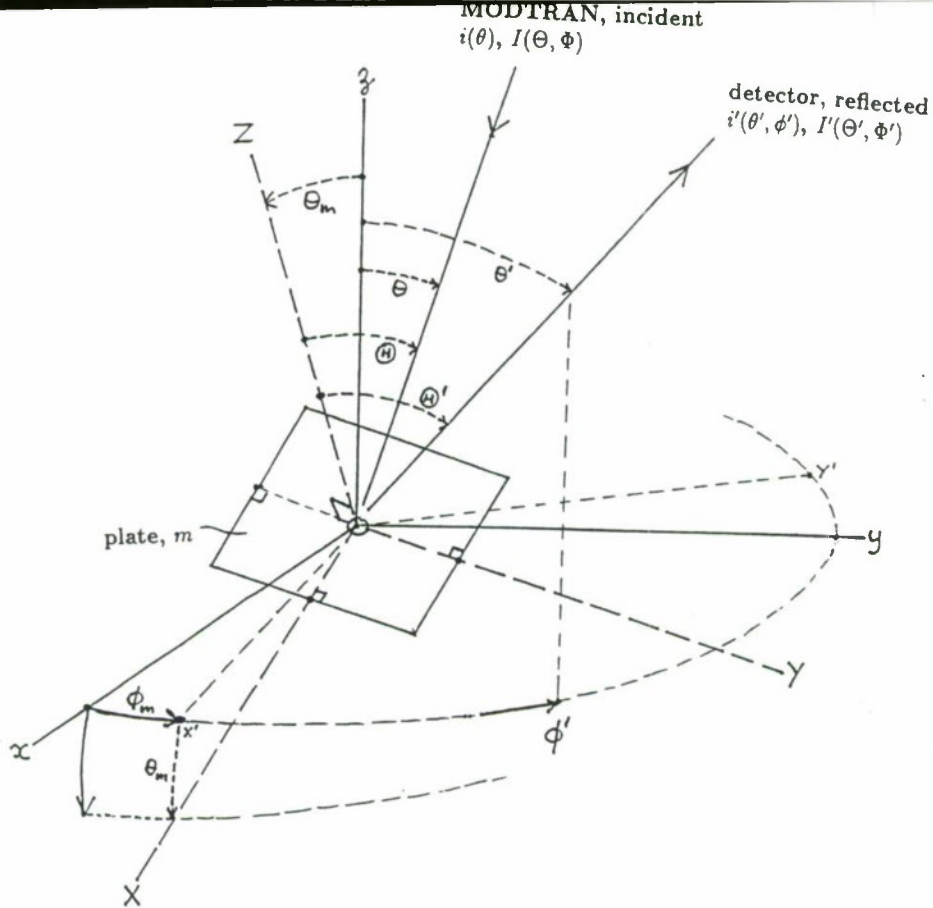
Fig.1 Orientation of plate- and earth-axes
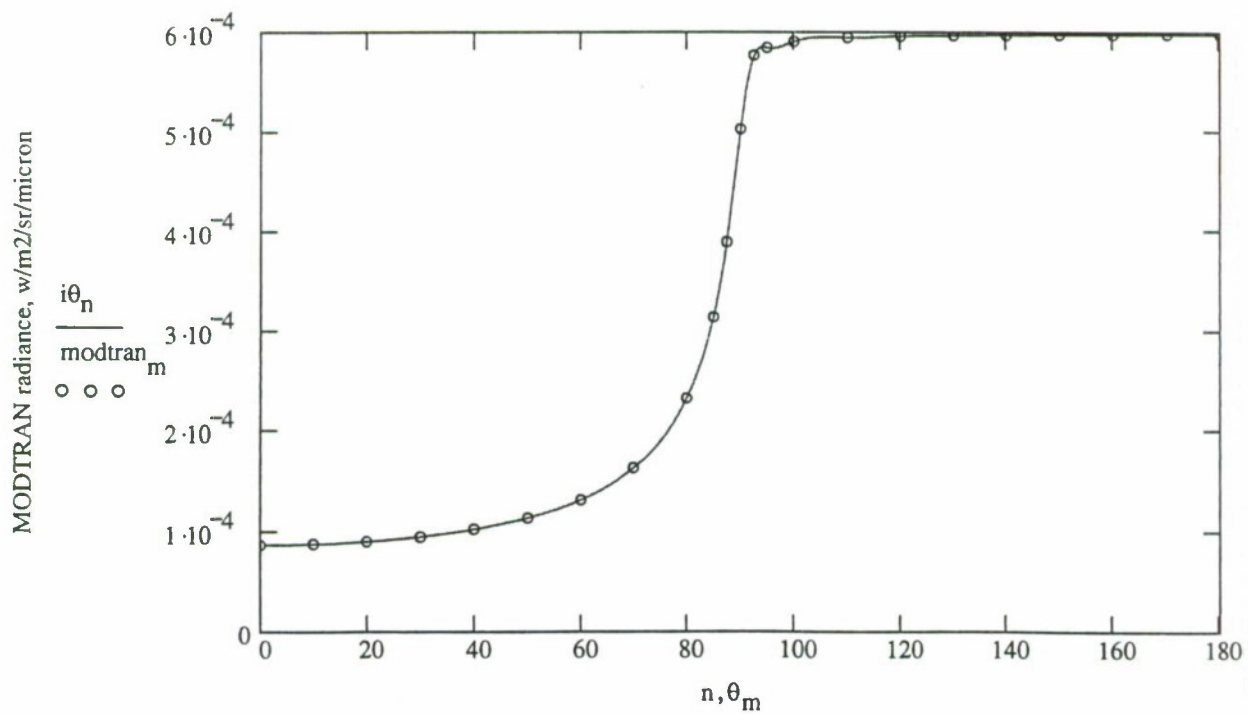
Fig.2 MODTRAN radiance $i(\theta)$ v. elevation angle θ

103
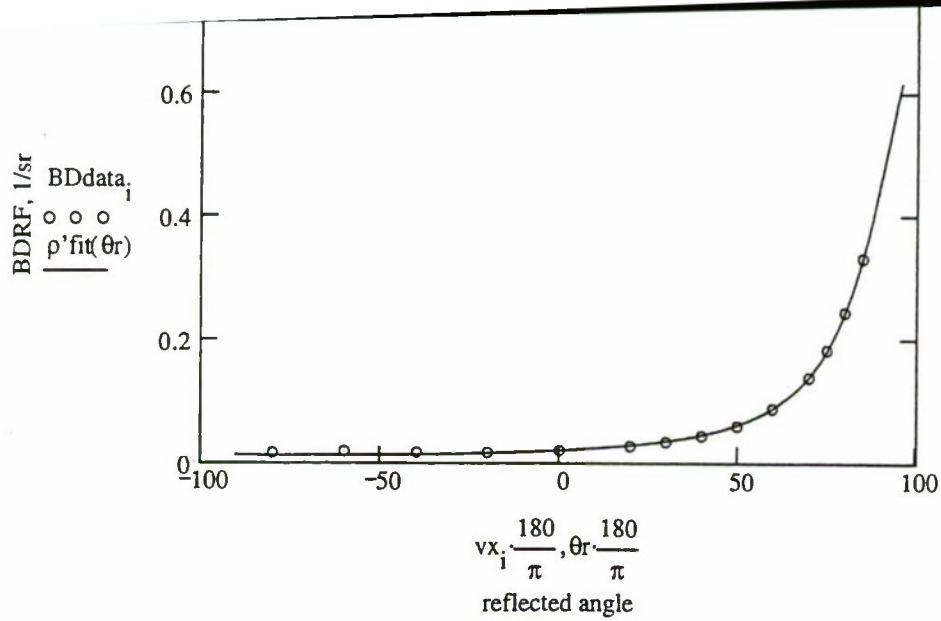
Fig.3 Fitting BDRF ρ' to measured data, $\theta_{inc} = 40°$
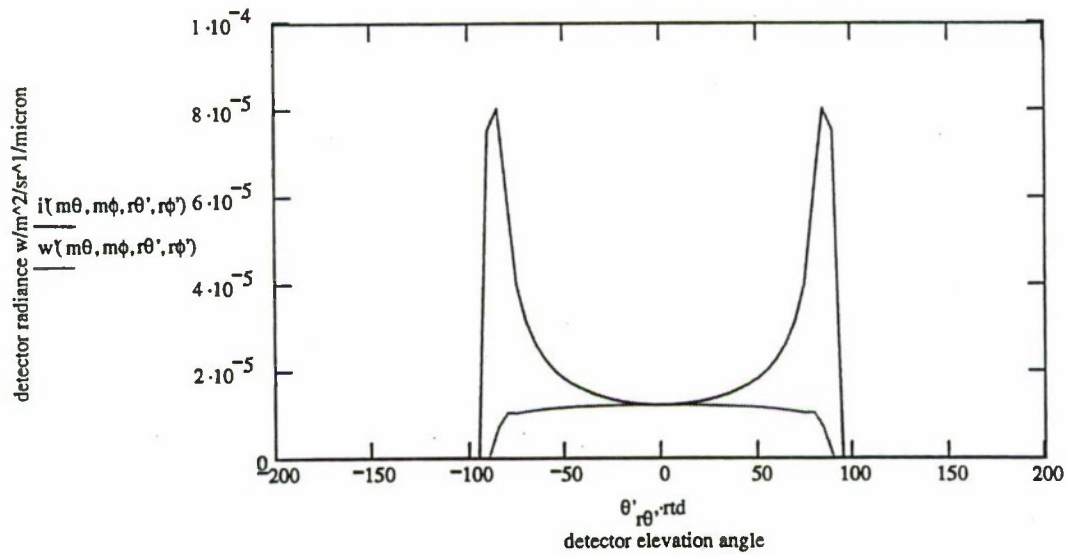
Fig.4 Detector radiance $i'(\theta')$ v. detector elevation θ', plate elevation $\theta_m = 0°$
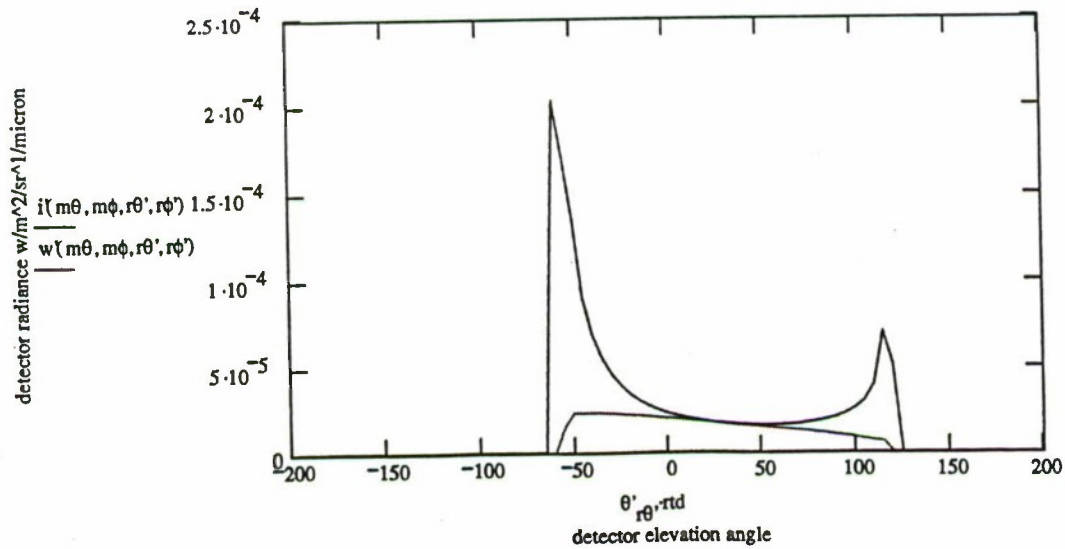
Fig.5 Detector radiance $i'(\theta')$ v. detector elevation θ', plate elevation $\theta_m = 30°$
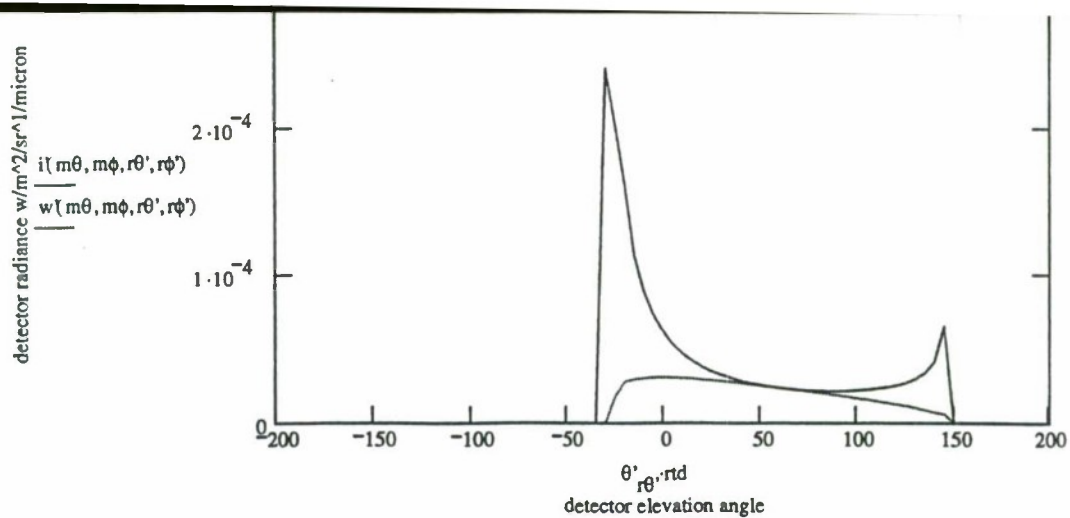
Fig.6 Detector radiance $i'(\theta')$ v. detector elevation θ', plate elevation $\theta_m = 60°$
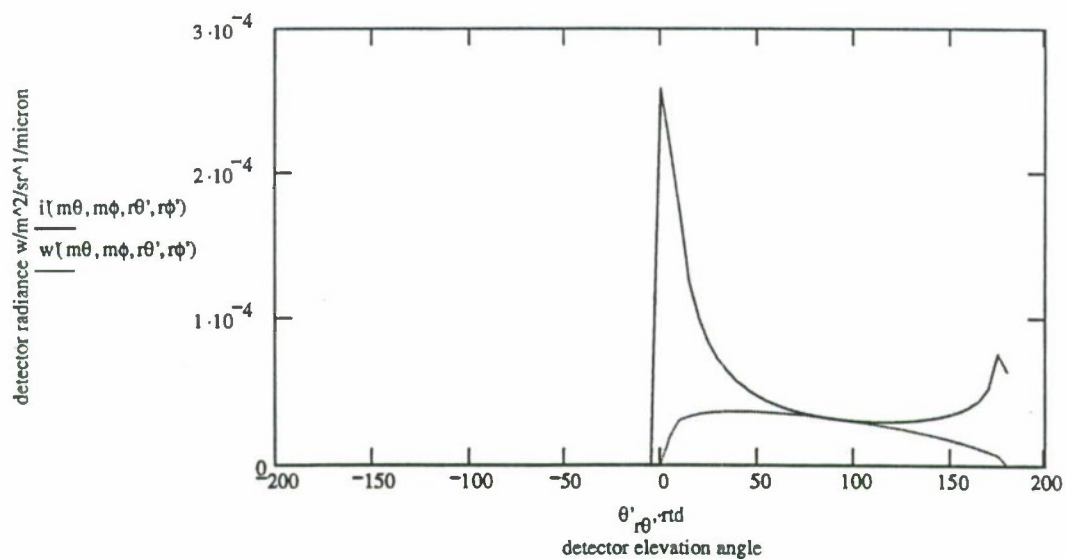
Fig.7 Detector radiance $i'(\theta')$ v. detector elevation θ', plate elevation $\theta_m = 90°$
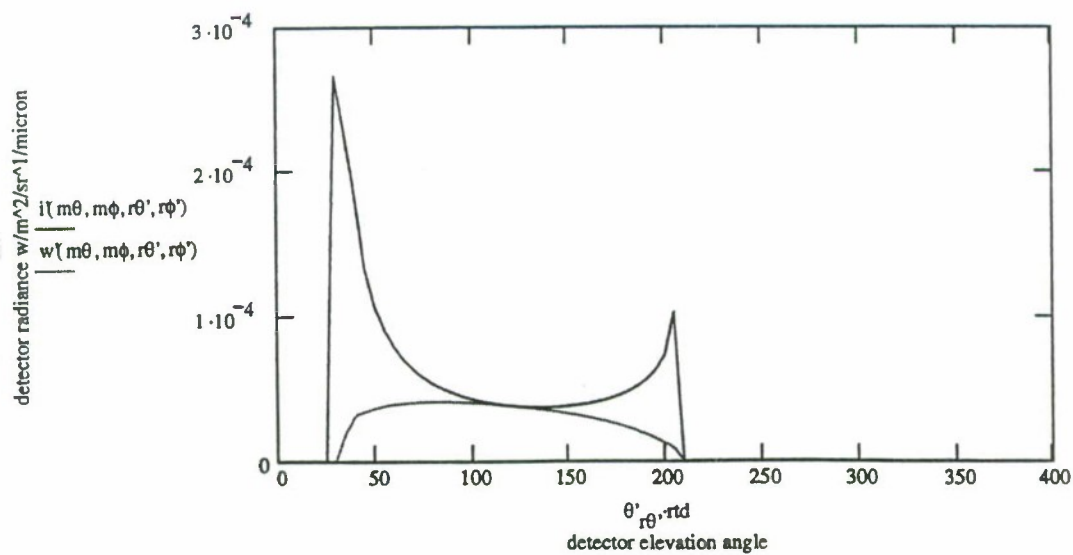
Fig.8 Detector radiance $i'(\theta')$ v. detector elevation θ', plate elevation $\theta_m = 120°$
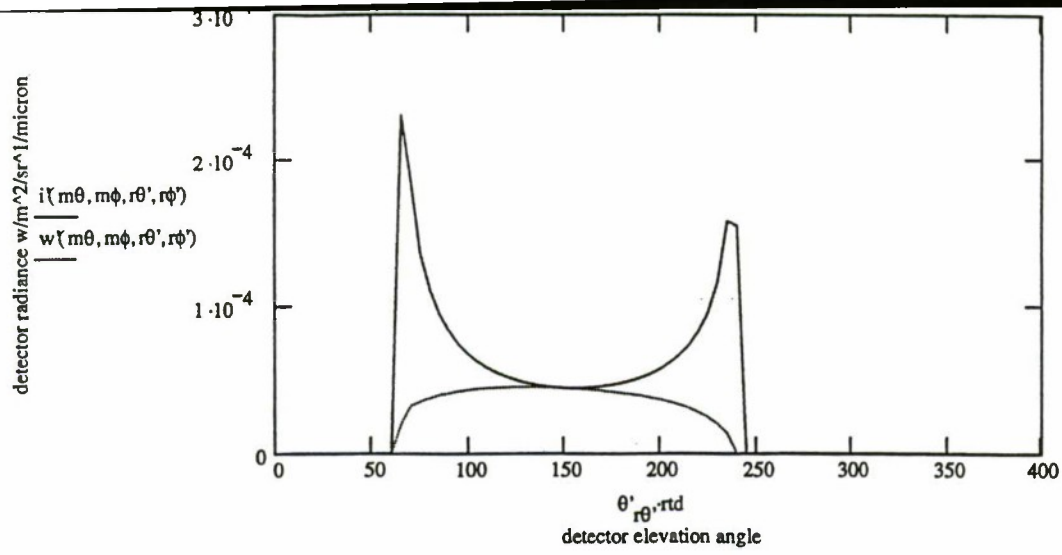
Fig.9 Detector radiance $i'(\theta')$ v. detector elevation θ', plate elevation $\theta_m = 150°$
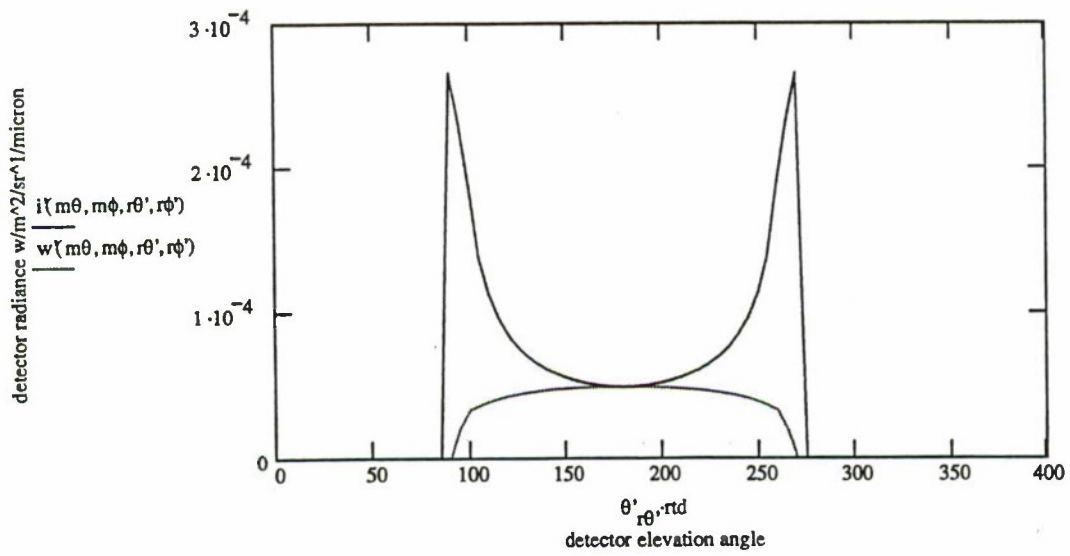
Fig.10 Detector radiance $i'(\theta')$ v. detector elevation θ', plate elevation $\theta_m = 180°$

106

BRDF Modeling for Physically Accurate Image Rendering

John W. Hilgers, William R. Reynolds, Randall J. Houle
Signature Research, Inc.

James Jafolla
Surface Optics Corporation

ABSTRACT

The bidirectional reflectance distribution function (BRDF) [1] describes the directional reflective properties of a surface. It is therefore a function of four angles, the azimuth and zenith angles for both the incident and reflected directions. If the reflectance is anisotropic, a function of the azimuth angle, then approximately 10^9 floating point values are required to specify the BRDF for one degree angular increments. In addition, if the surface has even moderately specular facets the BRDF will have lobes of large magnitude but small angular subtense. This makes utilization of the BRDF difficult from a practical standpoint.

A method for solving these problems was based on determining each lobe's location and angular subtense as a function of the incident direction angles. The lobe's shape was determined by computing its second moments and equating these to the corresponding moments of a Gaussian distribution function. Thus for each lobe six variables, the lobe's centroid, height and the three coefficients in the exponent of the Gaussian, were tabulated as functions of the incident direction. These functions were then modeled by bilinear interpolation. While straightforward conceptually, obtaining this database for highly structured surfaces proved a challenge.

Once the BRDF model is available it is applied to the problem of rendering physically accurate images. A unique hybrid rendering approach is described and demonstrated using anisotropic BRDFs. In summary the approach uses a measured background and employs first the radiosity method and then standard ray tracing to make extremely detailed renderings.

INTRODUCTION

The objective of the present study is to synthesize a calibrated spectral rendering of targets from the UV through the LWIR (8-12 microns) bands. Furthermore, arbitrary, realistic backgrounds are to be incorporated by projecting the imagery onto a hemisphere surrounding the scene.

This approach requires the rapid computation of many BRDF values, which can only be achieved by initially modeling these potentially pathological functions

accurately. The BRDF is obtained by experimental measurement, which results in an exceedingly large database, or, alternatively, the execution of a simulator like *MicroOpt*. It is assumed the BRDF is anisotropic and will eventually include polarization effects. The unique and innovative features of this approach include:

- the rendering of both diffuse and specular surfaces from a spatially differentiated background.
- The utilization of both radiosity and raytrace approaches to exploit the advantages of each.
- The insertion of the target into measured backgrounds and the addition of sensor effects.

FITTING THE BRDF

Due to the exceedingly pathological behavior exhibited by typical BRDF's, standard methodologies for the modeling and approximation of functions were impractical and alternate strategies were needed. Standard approaches [2] include interpolation, fitting a linear model or fitting a nonlinear model to the data. The first two methods usually involve the solution of a linear system; the third a Powell or Levenburg-Marquardt method for minimizing over model parameters. While such methods are very effective in single instances, applying them when thousands of incident directions must be evaluated is impractical. To circumvent these difficulties a method which obviates the need for optimization by computing approximations in the "forward direction" was applied. The method of choice was to fit a Gaussian distribution to the BRDF by equating moments.

Many different models were tested including wavelets, spherical harmonics, B-splines, Cauchy distributions, the OPTASM [4,5] model and the Gaussian distribution. Only the latter had finite second moments computable in closed form. By the Gaussian distribution is meant the function,

$$(1) \qquad \rho = Re^{(-ax^2 - 2bxy - cy^2)}$$

where $x = \theta_r - \bar{\theta}$ and $y = \phi_r - \bar{\phi}$. Equation (1) has six parameters. R gives the height of the lobe, the ordered pair $(\bar{\theta}, \bar{\phi})$ gives the lobe's position, and the triple (a,b,c) gives the lobe's orientation.

The zeroeth and second moments for (1) are:

(2)
$$\mu_0 = \iint e^{(-ax^2 - 2bxy - cy^2)} dxdy = \pi / D,$$
$$\mu_{xx} = \iint x^2 e^{(-ax^2 - 2bxy - cy^2)} dxdy = \pi c / 2D^3$$
$$\mu_{yy} = \iint y^2 e^{(-ax^2 - 2bxy - cy^2)} dxdy = \pi a / 2D^3$$
$$\mu_{xy} = \iint xy e^{(-ax^2 - 2bxy - cy^2)} dxdy = -\pi b / 2D^3$$

where in (2) $D = \sqrt{ac - b^2}$. It is easy to invert (2) which yields,

(3)
$$a = \mu_0 \mu_{yy} / D, \quad b = -\mu_0 \mu_{xy} / D$$
$$\text{and} \quad c = \mu_0 \mu_{xx} / D,$$

where this time in (3) $D = 2(\mu_{xx}\mu_{yy} - \mu_{xy}\mu_{xy})$

The procedure for the six parameters now is:

(4)
- Determine a reasonable sample frequency based on lobe behavior.
- At each sample point find and delimit all lobes.
- Then numerically perform the integrals in (2) on the lobe data.
- Then use (3) to establish parameters a, b and c.
- Let R in (1) be the lobe maximum.
- Let ($\bar{\theta}$, $\bar{\phi}$) be the lobe centroid computed from the first moments or, alternatively, it could be the (θ_r, ϕ_r) where the lobe maximum occurred.
- All six parameters are tabulated for each lobe.
- For given incident direction, interpolate these tables to approximate the six parameters.
- These parameters with the reflected angles allow the BRDF computation via (1).
- This must be done for each lobe, and a diffuse component added to fill in between lobes.

LOCATING AND DELIMITING LOBES

There are two general approaches for finding lobes at the sample points. Each has its advantages and disadvantages and each method is partially dependent on the other. Method 1 is the "tracking" program which, once a lobe is acquired, tracks it to the next sample point by gradually varying the incident direction. Method 2 is the "scanning" program which simply advances to the next sample point and searches for all lobes by computing on a very fine (usually angle increments of one or two degrees) reflection grid. Method 2 has the advantage of being automatic, albeit very time consuming. However, lobe identification at adjacent sample points still requires a rudimentary tracking program. Only by following the lobe can the user be sure which lobes at adjacent sample points are, in fact, the same lobe.

Method 1 is much less automatic since the tracking programs run the risk of losing the lobe for reasons cited below. The user may be required to run several different programs to reacquire the lobe, or, if the lobe has vanished, to determine if the lobe reappears. The tracking programs require a scan of at least one sample point to initially detect all lobes.

Some reasons why this part of the entire process is the most difficult and time consuming are:

- lobes vanish and reappear.
- Lobes coalesce and bifurcate.
- The hull enclosing the lobe may enclose multiple lobes.
- A tracking program slips off one lobe and onto another.
- Lobes can include the zenith which complicates anaylsis.
- Winding numbers values may disagree implying discontinuities.

These and other problems make post-processing of the lobe parameter databases, especially those for lobe height, R, and lobe position, essential.

SEPARATING DIFFUSE AND SPECULAR COMPONENTS

If the BRDF is computed by using equation (1) for the lobes exclusively, it will be orders of magnitude too small over regions far from any lobe. In the rendering process this can be disastrous since even small BRDF values can significantly impact pixel intensity in combination with strong sources.

This problem can be solved several different ways depending on the nature of the diffuse component of the BRDF.

The simplest fix consists in just adding a constant to the value obtained from the lobes using (1). This constant can be taken as the average of the diffuse component obtained by integrating over all four angles. Or the integration can be over the R-angles and the average value made a function of the I-angles.

The most accurate approach involves sampling the diffuse component coarsely and interpolating to get a value at each setting of all four angles.

If the more sophisticated methods are to be used it is necessary to separate the lobes from the diffuse substrate. This is most easily done by a filtering procedure. For example, if a tabulated or simulated BRDF value exceeds a tolerance, replace it with the average of its nearest neighbors. If necessary this process can be repeated until the lobes are sufficiently suppressed.

An alternate strategy for dealing with the diffuse component is to determine radiation exchange factors and compute the radiosity solution prior to executing the raytrace in the rendering code. This is discussed below.

Figures 1 and 3 show BRDF lobes for the Sanford-Robertson and Specular settings, respectively, in MicroOpt for two incident directions. Figures 2 and 4 show the reconstructions obtained by the method just described.
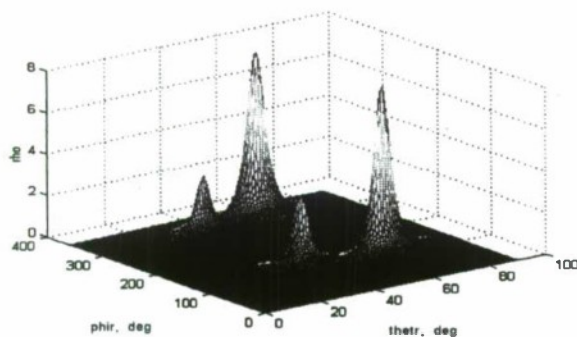
Figure 1: BRDF Sanford-Robertson option, I-dir = (40,0)

Figure 2: Gaussion Reconstruction for (40,0)
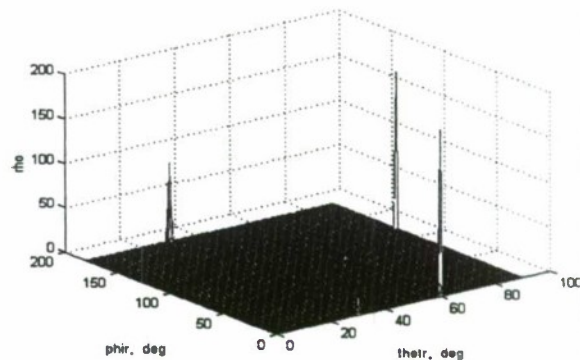
Figure 3: BRDF Specular option, I-dir = (30,0)
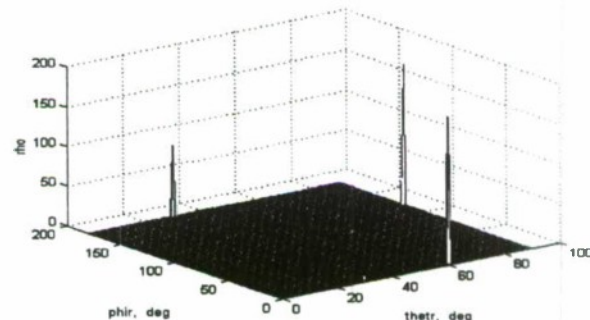
Figure 4: Gaussian Reconstruction for (30,0)

RENDERING APPROACH

Figure 1 schematically depicts the inputs and dependencies for the rendering process. BRDFs are modeled as described above and input to the renderer. FRED geometry models, RAD-X, which provides the radiation exchange factors for the radiosity solution, and PRISM are input in support of the rendering code. Currently these utilities remain "stand-alone", but eventually all components will be integrated into a single Visual C++ application.
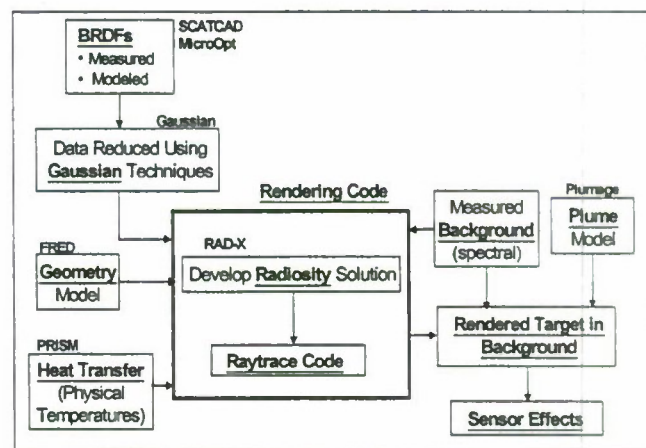
Figure5: Overview of Rendering Components.

109

The measured RGB background, as shown in figure 2, is projected onto a hemisphere surrounding the vehicle as shown in figure 3. Thus step 1, establishing the target radiation environment, is complete. Step2 is to obtain the radiosity solution. Figure 4 shows the radiosity solution for a hemispherical target with a single source.
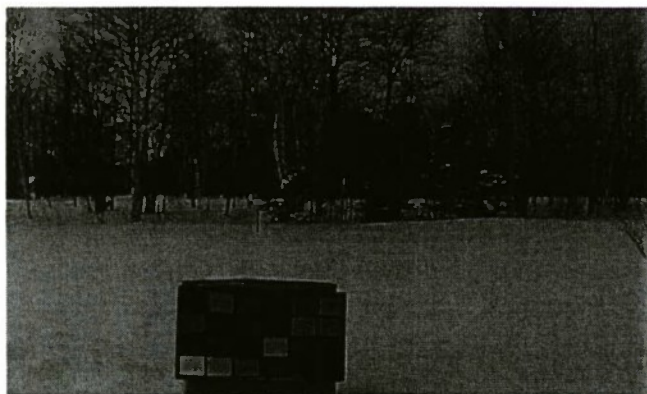
Figure 6: Image of Calibrated Background.

Figure 7: Projection of Background on Hemisphere.
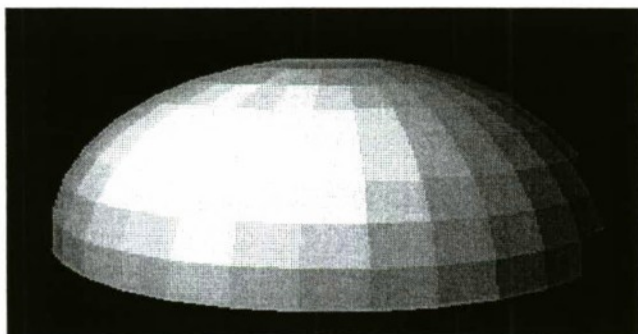
Figure 8: Radiosity Solution with Single Source.

Step 3 is to raytrace [3] along specular lobes and to defined sources as well as sufficiently strong intensities that result from the radiosity solution. Multiple specular reflections can also be implemented. Step4 is to create the rendered image by adding sensor effects. Figure 5 shows three panels

of varying specularity inside the hemisphere with the projected, measured background. The reflected tree line is clearly visible on the specular panel.

Figure 9: Three panels

SUMMARY

It has been shown that BRDF databases, whether measured or simulated, can be effectively modeled, even in extremely pathological instances that arise from highly specular coatings. The integration of the accurately modeled BRDF with ray tracing techniques has been shown to produce intuitively appealing rendered images, and initial validation runs confirm that this imagery is indeed physically accurate.

The continuing effort will be directed at making the BRDF modeling stage faster and more automatic. Better methods for dealing with inherent problems, like the behavior of the azimuth angle near the zenith, will be incorporated in the software.

In the rendering phase parameters that affect the acuity of the final image will be fine tuned. Greater accuracy will be sought by including multiple specular bounces. And eventually polarization effects can be included since the independent elements of the Meuller matrices can be modeled just as the BRDF can.

Finally the entire component package displayed in figure 5 is to be integrated into a coherent Visual C++ application.

Bibliography.
1. Thermal Radiation Heat Transfer R. Siegel and J.R. Howell (third edition) Hemisphere Publishing Corp. 1992.
2. Numerical Recipes in Fortran. W.H. Press et al (second edition) Cambridge U. Press 1992.
3. Practical Ray Tracing in C. C.A. Lindley John Wiley & Sons 1992.
4. Software User's Manual for the Optical Target ignatures Model (OPTASM). C. Acquista and R. D. Rosenwald PAR Gov. Sys. Corp. San Diego Technology Center, LaJolla CA.
5. OPTASM UPDATE. J. Stokes and J. Jafolla Surface Optics Corp. San Diego, CA.

Validation of Target and Terrain Models using Real-World Optical Properties Measurements

Jeffrey R. Kass
Surface Data Company
11555 Rancho Bernardo Road
San Diego, CA 92127-1441 USA

Chris B. Blasband
Surface Optics Corporation
11555 Rancho Bernardo Road
San Diego, CA 92127-1441 USA

ABSTRACT

Developing valid target and terrain modeling to use in night vision and infrared simulations is based on the liberal use of measured data from reliable sources. This data requirement remains the same, whether you model tactical scenes using high-resolution photographs, or larger areas from lower resolution satellite imagery. Industry partners Surface Optics Corp and Surface Data Company developed a system to produce quality spectral reflectance and thermal emissivity data. It covers the full gambit of field collection, laboratory measurement, and database production. The end result is a growing library of geospecific and geotypical data from .3-25μm, that makes results more believable and the output scenes much more realistic.

This paper describes the general approach taken to collect cultural and natural materials, measure their reflectance characteristics, and produce validated data for use in computer modeling and simulation. Maps, satellite imagery, aerial photographs, ground photographs, and geological literature are all used to decide what are candidate locations for collected materials for measurement and inclusion. A team is then sent to that area, if possible, to collect samples and document all activities. Samples are brought back to the lab for preparation and measurement. The data produced is then put into a general ASCII format, for easy extraction into any modeling system. This white paper shows sample grayscale images "before" and "after"- first with data from outside sources, then after applying data generated from our approach. A much higher level of structure is apparent.

INTRODUCTION

Using validated data is a two part process. Part one presents a description of the optical property data library produced by Surface Data Company (SDC) using Surface Optics Corporation (SOC) manufactured instruments and technicians. This library contains a continuously expanding set of surface materials and, for each, a data listing of its spectral reflectance from 0.3 to 25.0 μm (the actual data listings are given separately from this document).

Part two explains how applying this data supplies core validity to spectral modeling. This is shown through examining use of the optical data in a specific application.

Our approach uses "geospecific" data whenever possible. Information on geospecific samples are well documented and from identifiable locations. Samples that are of a more common nature, such as brick, yellow road stripe, and off-white exterior paint are considered "geotypical". Therefore, modeling the Yuma Army Proving Grounds, for instance, would have a much greater level of acceptance if it included optical properties of local "desert varnished rock", "Algodones Dunes sand", and "Arizona creosote bush" as opposed to generic granite, desert sand, and scrub brush.

The process for converting the rough samples into usable data follows an easy to duplicate process. This assumes that the data format is determined, based on the simulation in use. Under this system, data is reduced into seven major categories. However, the end user's application ultimately defines the layout of the data format.

THE OPTICAL PROPERTIES LIBRARY

The materials included in the SDC library are separated into seven categories. These are Soil, Rock, Vegetation, Hydrology, Construction, Fabrics, and Paints. We distinguish this material reflectance library from others available in that several types of reflectance data are provided, including reflectance as a function of the angles of incidence and reflection, and the diffuse and specular components of the total hemispherical reflectance. Libraries from other agencies typically provide only total hemispherical reflectance.

We've also added photographs with this library set to provide users an ability to visualize the characteristics of the samples and understand the environment from which they were extracted. The photos are scanned from photographs or digitally acquired.

The actual samples are compared to the video display on the computer monitor to adjust the color and tint to match the digital rendering as closely as possible. This is very important because the majority of visual imagery is not consistent with reality. This is especially true of false-colored "RGB" satellite imagery. The physical comparison immediately corrects this discrepancy.

Photographs are Included

Typically, three photos are included per sample (see Figure 1):

1. A general area shot provides a view of the local environment.

2. A site photo shows the immediate location from which the sample was obtained. Every effort was

3. A scanned image of the actual measured sample allows you to view the texture up close.

Naming Convention

The reflectance data files for each material are sorted by data type, with each file named using the following naming convention:

"Category_material name_data type_ID number"

For example:

"S_Sandstone_HDR20_AK0056" is the name for sandstone soil, hemispherical directional reflectance at 20 degrees incidence angle with ID number AK0056.

Geospecific Materials

Each naturally occurring earth surface material in the library (which includes soils, rocks, and vegetation) is specific to a selected geographical location. For these, photographs and their descriptions are given for the general area from which the sample was selected, the specific in-situ location of the sample, and the actual sample that was measured.

Also included are the actual latitude and longitude coordinates (WGS-82). In those instances where the same apparent material is included in the library, such as that of a dry lakebed or a dry silt-clay-sand from a desert wash, its specific location is the only information that can be used to explain differences in the reflectance of the two materials. This is most important for soils, and rocks whose mineral composition is difficult to determine, and

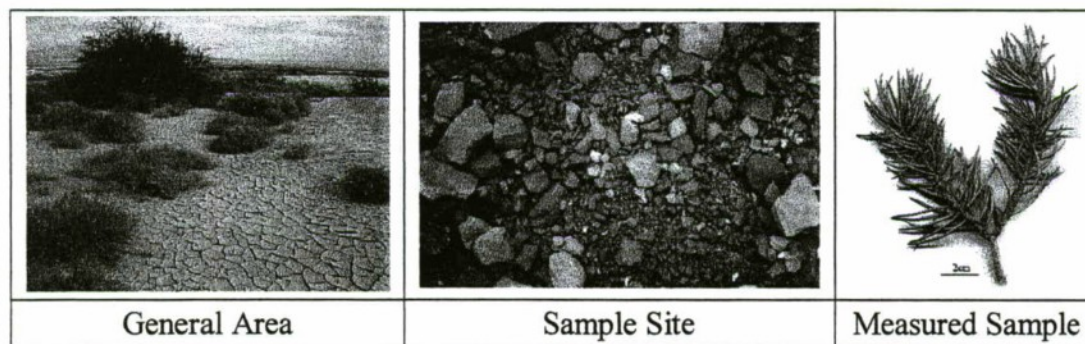| General Area | Sample Site | Measured Sample |

Figure 1, Sample photos included in library.

made to contain the field of view to 1-2 square meters.

least important for vegetation (for example, Honey Mesquite from Arizona has essentially the same reflectance as that from California).

The manufactured samples, in categories Construction, Fabrics, and Paints, are not specific to a location. For these, only a photograph of the measured sample and its description are given.

The Specific "Data Types"

Total Hemispherical Reflectance (THR) – Total reflectance for a near-normal angle of incidence (7 degrees) of illuminating source.

Hemispherical Directional Reflectance (HDR) – Hemispherical reflectance for specific angles of incidence of the illuminating source. The presently included angles are 10, 20, 40, and 60 degrees (these are denoted by HDR10, HDR20, HDR40, and HDR60, respectively).

Diffuse Directional Reflectance (DDR) – Hemispherical reflectance for specific angles of incidence of the illuminating source with the specular component blocked ($\pm\frac{1}{2}$ degree from the specular angle). The presently included angles are 20, 40, and 60 degrees (these are denoted by DDR20, DDR40, and DDR60, respectively).

Specular Directional Reflectance (SDR) – Specular reflectance component for specific angles of incidence of the illuminating source, computed by subtracting HDR and DDR. The presently included angles are 20, 40, and 60 degrees (these are denoted by SDR20, SDR40, and SDR60, respectively).

Bi-directional Directional Reflectance (BDR) – Reflectance as a function of reflectance angle in the plane of illuminating source for each incidence angle for which SDR data are measured. These data are provided in three spectral bands (visible, 3-5μm, and 8-12μm) from which the lobe width of the specular lobe can be determined.

FIELD WORK

Generating measured datasets starts with proper planning. First, the area of interest and targets for modeling are defined, based on an operational scenario or testing requirements. This drives the level of detail required and how many materials need to be represented at a given resolution. Both natural and cultural objects are considered. Sometimes data may already be available through our in-house library. Maps, satellite imagery, aerial photographs, ground photographs, and geological literature are all used to determine candidate collection locations.

Then, a team is sent to that area, to collect samples and document all activities. (Note: Sometimes an area is not

accessible, so a neighboring location with similar geology and vegetation is used.) Every effort is taken to comply with local and federal environmental restrictions and historical preservation policies. Samples and the general area are photographed, GPS coordinates obtained, prevailing conditions recorded, and locations are marked on maps.

In some cases, full diurnal thermal emissivity profiles of selected samples are captured on location. This is can be folded into the simulation, as required.

LAB WORK

Overview

Samples are brought back to the lab for preparation and measurement. After being logged into the database, SOC technicians use several instruments including the SOC-100 (Figure 2), SOC-250, and a CARY for the .3-2.0μm range. All instruments are run against a "baseline

Figure 2, SOC-100 Instrument with grass sample about to be measured.

standard" prior to the first sample measurement to ensure consistency. (Note: SOC uses "fused silica".) Experienced engineers and technicians scrutinize the results of measurements for data as part of the quality assurance process.

The data produced is then put into a general ASCII tabular format, for easy extraction into any modeling system. Figure 3 shows a plot of sample soil data
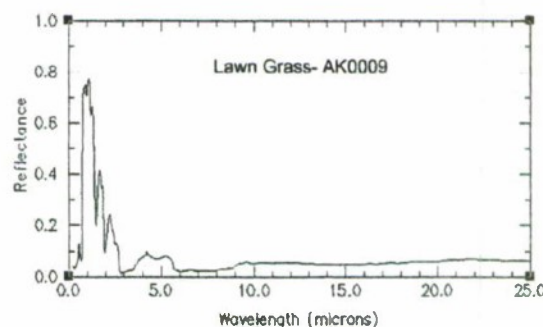
Figure 3, Reflectance plot of grass sample.

presented as reflectance, as a function of wavelength.

The area of interest can now be satisfactorily modeled based on data representing materials actually present. Although simulation algorithms require their own evaluation as to simulation accuracy, the underlying data is considered reliable. Auditing of data generated at this core level is easy, since all samples, except vegetation (new samples can be obtained based on GPS coordinates), are retained for re-measurement, upon request.

Check List

The following checklist is used in working with end users under validation and verification (V&V) program requirements.

- ❑ Tasking formally defined, with designated area of interest to model
  - o Defined level of detail required and resolution needed
  - o Defined V&V level of effort
  - o Budget and schedule
  - o Intelligence resources are available

- ❑ Candidate samples chosen from SDC library versus unique sample requirements

- ❑ Field measurements versus lab measurements
  - o Field and Lab instruments are identified, based type of samples listed
  - o Data recorded and documented

- ❑ Data processed
  - o Reduced data to user defined format
  - o Packed for delivery

CASE STUDY

Geospecific data generated from a field trip to the Death Valley area was used in the creation of radiance texture. This was done to compare these results with a previous rendering, using best available geotypical optical information. The end result was an impressive improvement in both the overall radiance information produced and in the structure of the infrared texture generated.

Software Packages Used

Two commercial off the shelf (COTS) modeling tools were used in this case study:

- ▪ Texture Material Mapper™ *[1] (TMM). This is a software tool that allows a user to associate optical properties to "RGB" pixels in a texture image.

- ▪ MOSART Atmospheric Tool™ (MAT). This is a software tool that generates the underlying atmospheric and natural thermal data for sensor simulations.

We started by importing a false color, LandSat RGB texture of Panamint Valley, CA[2] into TMM and applied the optimal geotypical materials, which were originally delivered. Figure 4 shows this list.

Selected Materials:

SOIL.limestone_rock
SOIL.salt-silt
SOIL.desert_sand
SOIL.lake_sand
SOIL.road_gravel
CONSTRUCTION.black_asphalt
SOIL.varnished_sandstone
VEGETATION.dry_grass
VEGETATION.pine-scrub
VEGETATION.lawn_grass
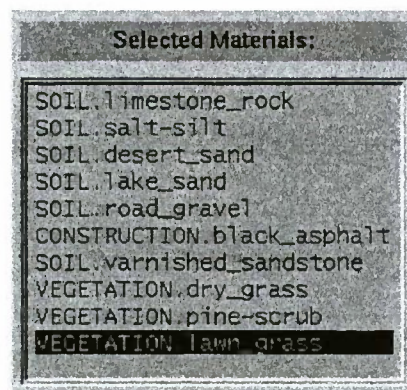
Figure 4, TMM geotypical data used in this demonstration.

Then, still using this tool, we generated a radiance texture using the parameters as shown in Figure 5.

---

[1] TMM & MAT are property of MultiGen-Paradigm, Inc.
[2] Provided courtesy of Raytheon Systems, Inc.

114
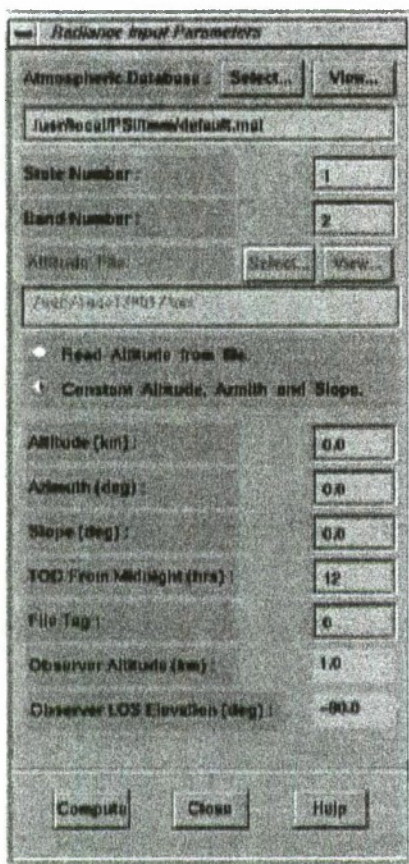
Figure 5, Radiance texture settings.

A radiance texture was created (Figure 6) having the following values:

- Min radiance= 0.00031 w/cm2•sr
- Max radiance= 0.00090 w/cm2•sr
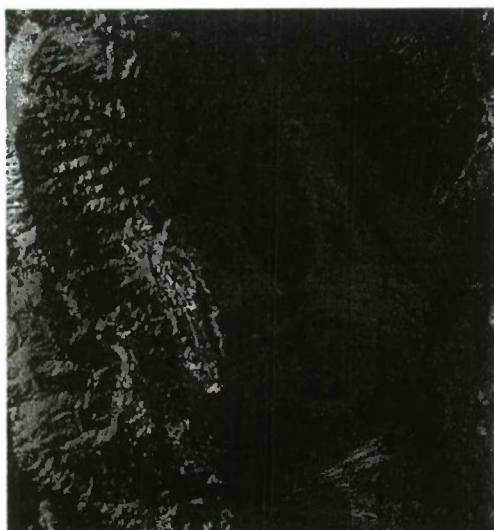
Figure 6, Radiance Texture from geotypical data.

Next we repeated the process, using some of the 77 new SDC geospecific materials (See Figure 7). The result

Selected Materials:

R.Breccia_THR_AK0035
R.Dolerite_THR_AK0030
R.Gabbro_THR_AK0092
R.Granite_THR_AK0070
R.Limestone-Caliche_THR_
R.Limestone_THR_AK0060
R.RockVarnish_THR_AK0073
S.SandDune_THR_AK0062
S.SandDune_THR_AK0071
V.BigSagebrush_THR_AK000
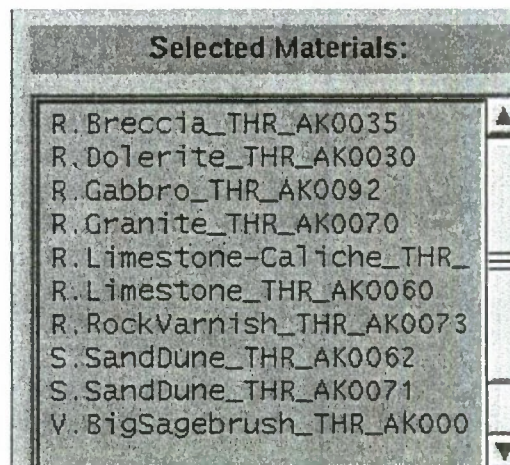
Figure 7, New geospecific materials applied.

from this application of data was:

- Min radiance= 0.00029 w/cm2•sr
- Max radiance= 0.00060 w/cm2•sr

Here is the radiance texture produced (Figure 8).

Figure 8, New geospecific radiance texture.

Comparing the Results

A side-by-side comparison reveals obvious differences in detail and image radiance. Here are some of those differences:

115

- Radiance contrasts of the two images are very different, particularly in lower right quadrant.
- Pixel-to-pixel radiance differences are as great as 50%.
- Geospecific maximum radiance is less by 33% and at different image location.

SUMMARY

Presenting results of modeling efforts to customers, or end-users, leaves a lasting impact on them. Incorporating verifiable, underlying data with a validated visualization system, promotes a high acceptance level. This type of data must be well documented. The validation process should then be quick and painless.

However, if the simulation uses questionable data, even though the visualization system is validated, results will suffer greatly. Additionally, accuracy of any follow-on results presented will always be suspect.

The following steps are required in geospecific modeling:

1. Perform a good field survey.
2. Use geospecific optical properties when appropriate. Make sure it is validated.
3. Compare results with existing imagery, when possible.

REFERENCES

[1] Champetier, R. J. & Friese, G. J., Aug 2, 1974, "Use of Polished Fused Silica to Standardize Directional Polarized Emittance and Reflectance Measurements in the Infrared," Report #TR-0075 (5450-60)-1, Air Force Systems Command, LAAFS, CA.

[2] Cornette, W. M., Acharya, P. K., & Anderson, G. P., "Using the MOSART Code for Atmospheric Correction," Aug 1994, Paper #IGARSS/1/1//06490, IEEE International Geoscience and Remote Sensing Symposium.

[3] Anding, D. C., & Szabo, A., "Real-Time Image Visualization for Sensors," Apr 1996 SPIE Conference, Orlando, FL.

SIGNATURE SITUATIONAL AWARENESS
FOR SURVIVABILITY ENHANCEMENT

Donald W. Rees, Otto Renius
General Dynamics Land Systems
Sterling Heights, Michigan 48310

William R. Reynolds, H. Griffis
Signature Research, Inc.
Calumet, Michigan 49913

ABSTRACT

General Dynamics and Signature Research have been developing a signature based situation awareness model to generate a maximum survivability route plan for ground vehicles. The model, called the Signature Situational Awareness Aid, (SSAA) incorporates user input for definition of parameters such as weather, environment and operating conditions together with a stored database on vehicle systems, terrain maps, threat characteristics and threat location to calculate multi-spectral signatures, line of sight information and detection vulnerability. The output generated includes a situational map, vulnerability status and probability of detection in both static and dynamic modes. Plans are in progress to incorporate all output information into the generation of a route plan for the vehicle to follow to its objective which minimizes its probability of detection from known and postulated autonomous and crew served threat weapon systems.

INTRODUCTION

The Joint Chiefs of Staff's vision of future warfighting stresses the development and use of information superiority to achieve full spectrum operational dominance.[1] A major factor in achieving and using this information superiority is obtaining a real time dynamic situational awareness capability to aid in decision making. As defined by the DoD, situational awareness is the "knowledge of one's location, the location of friendly and hostile forces, and external factors such as terrain and weather that may affect one's capability to perform a mission"[2] The Signature Situational Awareness Aid being developed by General Dynamics and Signature Research addresses the information superiority objective for ground vehicles through a signature based model to generate a minimum detection vulnerability route plan using both dynamic sensor input and stored data. Feedback to a command and control grid is provided for continuous update of battlespace awareness, and a real time input of the tactical situation is sent to the precision force to present current information on friendly and enemy force location and tactics.

Figure 1 illustrates the information flow using signature situational awareness as one facet of achieving full spectrum operational dominance.
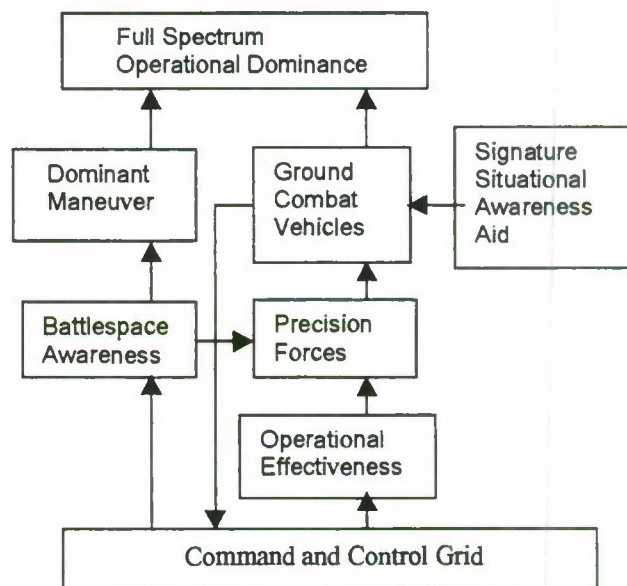
Figure 1 Information Flow to Achieve
Full Spectrum Operational Dominance

Signature Situational Awareness Aid Methodology

The SSAA was conceived as a way to enhance combat vehicle survivability by automating the input of multi-spectral signature data into a situational map to generate a maximum survivability route plan to the vehicle's objective. Figure 2 illustrates the overall methodology for the SSAA. In all cases, the input can be provided from stored data, a command center, or entered directly by a member of the vehicle's crew. Calculations are performed using software specific to SSAA, but which incorporates look up data as required. Output displays are provided to the vehicle crew to help in decision making, and information can also be sent to a command center to add to battlespace awareness as an aid in decision making and an update on the tactical situation.
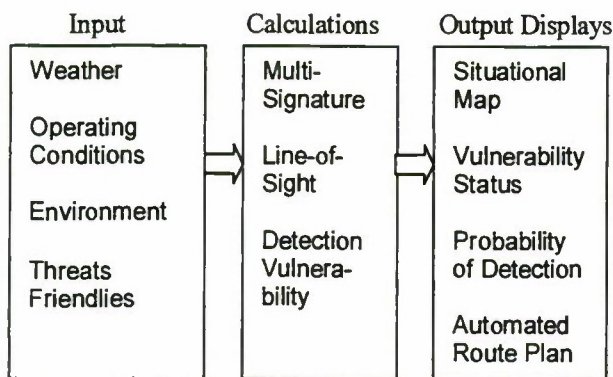
Figure 2 SSAA Methodology

Input

The input to SSAA is a combination of sensor observations, data from the command grid, and crew observations. For example, weather data can be entered by the crew through observation. However, as the SSAA become more mature, it may become desirable to incorporate sensor data such as solar irradiance history to provide a more accurate self signature calculation, and relieve the crew from the burden of entering this data.

In its current form, SSAA weather input is limited to relative humidity, cloud cover, rain rate, and visibility distance. The input can be entered from meteorological information at the beginning of a mission, and updated by the vehicle crew as conditions change. Figure 3 illustrates the weather input selections as they appear on the SSAA screen.

Background selections are made to describe the vehicle's immediate area. In the SSAA's current configuration, the background is divided into organic materials such as grass, weeds, shrubs, and trees, and inorganic materials such as sand and concrete. The vehicle crew can update the background data as the background varies during dynamic operation. Figure 4 illustrates the background screen input to SSES.

Figure 3 Weather Input Selection
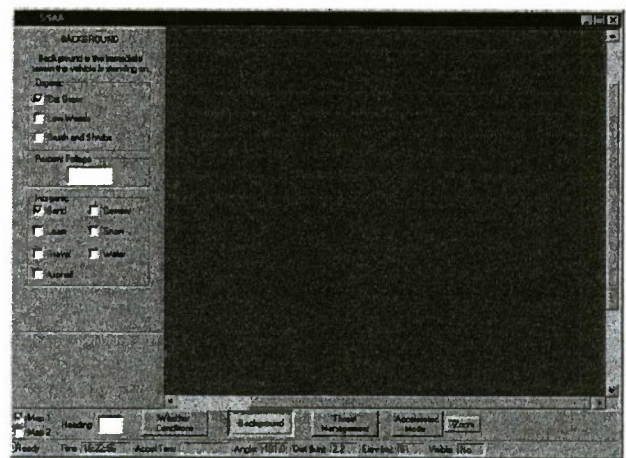
Figure 4 Background Parameters

Both the friendly (blue) force and threat (red) force type and location are entered into the SSAA. Currently, the planned blue force input to SSAA is limited to the M1 tank, M2 IFV, M-113 APC, and HMMWV truck. Sensor characteristics of each vehicle are stored in the computer's database. The red force location and makeup is also entered into the SSAA, and displayed on the SSAA's screen. Stored data on each type of vehicle, both blue and red, is used to calculate the vehicle's capabilities. For example, the T-80's FLIR detection capabilities are available for determining blue force detection vulnerability. The force location and mix can be altered as information from external sources is received. Figure 5 illustrates a portion of the red force selection menu. The stored data on each threat system includes the type and performance characteristics of the target detection equipment (FLIR, I^2, etc.) that is normally mounted on that vehicle. This allows the calculation of the threat search and detection capabilities.
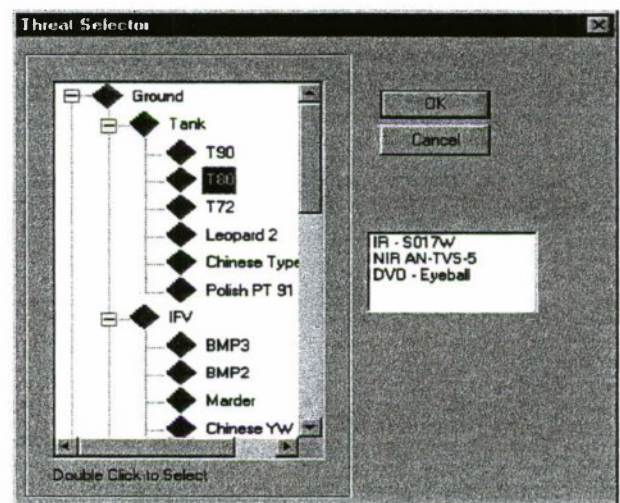
Figure 5 Red Force Selection Menu

118

As stated in the introduction, situational awareness for full spectrum operational dominance requires knowledge of the friendly force location. This is accomplished in the SSAA by accommodating the input of the type and location of all known friendly vehicles. Transmitted data from the central command and control grid allows the precise location of the vehicles to be displayed on the SSAA screen. Line of sight and detection probabilities for the friendlies are not included in further calculations, but identification of other blue vehicles provides the vehicle crew with updated battle-space awareness in their operational area.

Vehicle Operating Conditions

Current production U.S. combat vehicles are equipped with a digital electronic architecture that includes sensors, data busses, microprocessors and memory. This provides the opportunity for sensing, displaying and storing a time history of the vehicle operating conditions that contribute to the calculation of vehicle – background thermal, visual and radar signature contrast. In addition to main engine operating data such as time spent at idle RPM versus higher RPM movement at speed, operation of auxiliary power generators, if incorporated on the vehicle, provide a time history which must be considered in accurate thermal signature calculations.

Most vehicles are also equipped with on-board navigation systems that provide information on their current latitude, longitude, elevation, and hull heading direction. A turret position indicator also shows the turret/gun heading with respect to the hull. These systems are capable of providing the time history data necessary for determination of relative sun angles for various vehicle surfaces facing the threat and their exposure to solar load. These exposed vehicle surface data are also used in calculating the effective radar cross section and visual vehicle/background contrast due to radar and solar illumination of the various vehicle surfaces. In addition, the navigation and turret position indicator systems also provide information on which side of the vehicle is facing the threat and the direction of the gun relative to the threat.

Calculations

Once the input information on weather, environment, operating conditions and threat is entered, the SSAA calculations can be made to estimate the blue vehicle's current vulnerability to detection by various threats.

At present, predicted detections by threat sensors in the IR and visible bands are incorporated in the SSAA. Future efforts will implement threat acquisition estimates in radar, MMW, and acoustic bands. For the RSTA and EO sensors presently implemented in the SSAA the path to obtaining a probability of detection, P_d, takes the in-band self signature, propagates it through the atmosphere at range and then uses characteristics of the threat sensor to compute a P_d.

The signatures of blue target vehicle along with its surrounding background must be estimated prior to calculating the probability of detection by a threat sensor. In general, to estimate the threat P_d, requires knowledge of the threat sensor performance, range to the threat, atmospheric transmissivity, and the vehicle signature including aspect and background. Although having limitations, we have incorporated the ACQUIRE model formalism to generate the actual P_d estimates. In general, the ACQUIRE Model works well for resolved targets having moderate contrast in backgrounds with medium clutter levels.

Line of Sight Determination

Standard military topographic maps available on CD-ROM are utilized in the SSAA to determine line of sight (LOS) between the blue and red vehicles. SSAA elevation data was MIL-PRF-89020A standard Digital Terrain Elevation Data (DTED) Level 1 with three arc second accuracy. The 1:125000 scale SSAA topo map is generated from the DTED Level 1 data with elevation gradients of 20 meters.

The LOS algorithm is based on the routine used on NIMA MUSW 2.1 software. In this algorithm, the depression angle between the observer and each point along the line to the target is computed. The angle is compared to the previous angles to determine if LOS is maintained.

IR Signature Estimation

The IR signature of the blue vehicle can be obtained by several different means including utilizing surface temperature sensors, look-back sensors, nearest neighbor sensors and thermal models. While the first three have definite advantages and could be implemented in the SSAA we have made the less intrusive choice of using the thermal model PRISM.

PRISM is a proven high fidelity, first principles, IR signature model having a long developmental history and has been widely reported in the literature. In summary, a PRISM signature is computed using a coarse nodal description of the vehicle where a dynamic energy balance is computed for each node. From this energy balance a node temperature is computed. Given the surface optical properties of each node or facet the radiosity of each surface is computed.

PRISM is unique in that the internal thermal sources are developed from first principles. For example, the power required at the drive sprockets is the product of the speed and rolling resistance plus the power required to negotiate the terrain grade. Therefore, power developed at the drive shaft is a sum of the transmission and gearbox losses and the drive sprocket power. The amount of waste heat is derived from the sum of the thermal losses due to inefficiency of the engine and transmission. The waste heat leaves the system via the exhaust gases and the cooling system. Given mass flow rates of the exhaust gases and coolant air the temperature of these gases can be accurately estimated. Consequently, the temperatures of the exhaust stacks and exit grill can be computed.

Some of the required inputs to PRISM are readily available. For example, vehicle latitude, longitude, elevation, speed and direction are available from on-board GPS via the vehicle vetronics bus. Asymmetrical solar heating of the vehicle is computed in PRISM from the heading information. At present the meteorological data as well as background characteristics are primarily obtained from user input over the diurnal cycle. If the SSAA is integrated into a future vehicle inexpensive on-board sensors could be used to provide higher fidelity data. Background information can also be derived from a high fidelity DTED database if available and subsequently can be supplemented by user input.

Visible Signature Estimation

We have implemented an empirical model of the luminance contrast of baseline military targets in back-country terrain that was derived by Gary Witus of Turing Associates. At present, first principles visual contrast models are not widely available or validated in general. The Witus empirical model is restricted to military ground vehicles with standard forest camouflage paint, in terrain with fields of dry grass and deciduous trees and bushes with green leaves. At present, the model does NOT account for any range-dependent atmospheric effects.

The model automatically calculates the luminance of the designated surround and three regions of the target signature:

1. The lower glacis (or front grill below the hood) and its shadow (if any),
2. The side of the vehicle and its shadow (if any), and
3. The upper glacis, upper deck and roof or top of the vehicle.

The model automatically calculates the contrast ratio of each of these regions relative to the surroundings.

The final output, the average contrast ratio, is the weighted average of the contrast ratios of the three regions, weighted by their proportion of the total presented area.

Visual Contrast Model Derivation

The model is based on measurements made from the 44 Search_2 high-resolution digital images. The image set contains 9 types of military vehicles (US and Soviet, tracked and wheeled, thin-skinned and armored), in a variety of locations and postures at ranges from 500 to 5,000 meters. Images were taken with a 10X lens, and the 35 mm slides were digitized at 6K by 4K pixels. The atmosphere was clear. In six of the 44 scenes, the target was in the shadow of a cloud. Dr. Lex Toet of the TNO Research Institute in the Netherlands provided the digital images and data to convert gray-scale to luminance values.

These scenes were staged to be representative of military targets in field positions. It is very important to observe that the targets had a very non-uniform appearance. Typically the upper glacis and top of the vehicle were very much brighter than the lower glacis and side, and much brighter than the tree and bush foliage. The lower glacis and side (if self shadowed) and shadow on the ground were typically darker than the tree foliage. Directly illumined sides were much brighter than self-shadowed areas, but typically darker than the roof or upper glacis. For these reasons, the average luminance over the target is not a sensible basis for a target signature metric.

The simple alternative approach taken was to divide the vehicle signature into three regions, compute the luminance and relative area of each of these three regions, as a function of illumination angle and viewing angle. For each image, the average gray-scale value was measured for the following three regions:

1. The front (rear) lower glacis or grill, front tracks/tires, its shadow, and frontal gun/turret shadow (if any);
2. The side of the chassis and turret (or hull and cab), side tracks/tires, and their shadow if any; and
3. The top roof, upper deck, and upper glacis.

The reason for selecting these regions lies in the geometry of the vehicle. The regions were (mostly) contiguous, and had about the same apparent luminance. This organization explained most of the variation in target pixel luminance. The average within-region gray-scale standard deviation was 10 (out of 255), whereas the average gray-scale standard deviation of the overall vehicle luminance was 35 (out of 255).

The average gray-scale values for directly illumined grass, shadowed grass, tree/bush foliage, and local target surround were also measured. For each image, the target aspect, lookdown angle, illumination angle relative to the target orientation to the observer, and whether or not the vehicle was in a cloud shadow were recorded.

From these data, the proportion of target area represented by each of the three regions, as a function of lookdown angle and target aspect were computed. The average gray scale value for each of the three regions was computed for each of the 4 illumination angle conditions (note that since the illumination angle designation is relative to the observer viewing angle, there is no need for a matrix of illumination angles and target aspect angles).

The model uses the gray-scale to luminance equation and constants contained in the TNO report "A high-resolution image data set for testing search and detection models," TM-98-A020, to calculate the average luminance of the designated surround and each of the three target regions.

The contrast ratio for each target region is simply the absolute value of the target region luminance minus the surround luminance, divided by the surround luminance. The average target contrast is simply the weighted sum of the contrast of each of the three regions. The weights are simply the proportion of the presented target area represented by each region.

Note that this is not necessarily the best or right way to measure contrast. It is the approach that most closely resembles the standard Army contrast ratio formula (target minus background, divided by background), extended to account for the fact that the targets consist of distinct regions with very different luminance. If the target had a constant luminance, this would reduce to the standard Army contrast ratio form.

Once the vehicle signature is determined, for the operating, environmental, and atmospheric conditions of interest, the probability of detection is calculated. Several additional data items are required for these calculations.

These data items include:

Threat Sensor Performance Data
Range from Sensor to Target
Atmospheric Effects
Terrain Data to Calculate Line of Sight (LOS)

The probability of detection calculations begin with a check to determine if the threat sensor has line of sight with the target system. An element of error in the location of the threat is introduced to determine if a threat system

within 1300 meters of the of the user determined threat location has line of sight with the target. If the threat has LOS then the probability of detection calculations are continued. If not, then the probability of detection is zero for that particular threat location. Once LOS has been determined, the probability of detection is calculated. The ACQUIRE methodology is used for man-in-the-loop sensors to determine the number of resolvable cycles on target. For man-out-of-the-loop sensors, other methodologies to determine detectability will be used.

Acquire

The ACQUIRE methodology uses vehicle presented area, atmospheric effects, vehicle signature, and the threat sensor performance (minimum resolvable contrast vs. spatial frequency curve) to determine the number of resolvable cycles on target. After the number of resolvable cycles on target has been calculated, an "N50" value is incorporated to determine the probability of detection. The "N50" value is the number of resolvable cycles on target necessary to detect the target 50% of the time.

The ACQIRE calculation yields a probability of detection for a specific sensor-target pairing under specific environmental and atmospheric conditions, and for specific vehicle operating conditions. This probability is displayed to the user by the SSAA in both a digital and graphical manner.
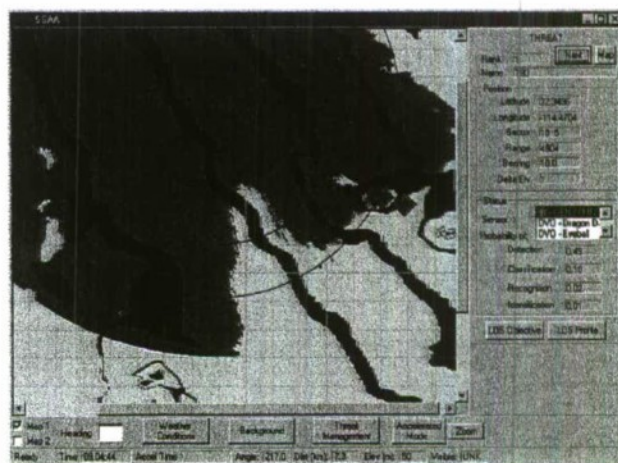
Figure 6   Clear Line of Sight and Detection

Figure 6 illustrates a typical SSAA display . In this case, the blue vehicle is represented by an Abrams and the red vehicle by a Russian T-90 tank. In this case, the probability of detection of the Abrams by the T-90's Generation 3 FLIR is 0.49 at a range of 4904 meters.

Figure 7   No Line of Sight

Figure 7 illustrates that under the same conditions, the Abrams would not be detected by the T-90 at bearing 75deg. and range 6255 meters because there is no line of sight between the vehicles. As the Abrams traverses the field, the probability of detection varies and is displayed on the SSAA screen for each threat selected.

Future Plans

Future SSAA enhancements will include a route planning capability as part of the overall mission planning mode. The SSAA route planning will be unique in that the acquisition probabilities of the threats will be incorporated with the standard route planning parameters such a line-of-sight to the threat, distance, terrain, etc.

If there are several routes that are possible and each exposes the vehicle to acquisition for some amount of time, t, and each at a different range then one can imagine trying to select between a route having a lower P_d for a longer exposure time versus a route yielding a higher P_d for a shorter exposure time.   This yields an exposure metric which is the integration over the total exposure time of the product $P_d(t) \cdot dt$ on the route. Other factors which must be considered are how the vehicle cross track speed impacts Pd because clearly the total exposure time could be minimized by an increase in vehicle speed where possible and potential search times for the field of regard.

In summary, the SSAA provides the warfighter with a powerful and unique route planning tool which incorporates the probability of acquisition by threat sensors, and will contribute to the US Army's goal of information superiority.

REFERENCES

(1) Joint Vision 2010 Chairman Joint Chiefs of Staff
(2) GAO Report 1998

UGV Modeling and Simulation in a SIL Environment

Rich Goetz and Grant Gerhart
US Army Tank-automotive RDE Center (TARDEC)
Warren, MI 48397-5000
goetzr@tacom.army.mil
gerhartg@tacom.army.mil

ABSTRACT

There are a variety of efforts attempting to integrate unmanned vehicle systems into the Army array of tools and weaponry. Military planners for Force XXI and Army After Next are busy assessing the viability and potential usability of unmanned vehicles of all types – land, sea, and air -- in the near, mid and far term. But exactly how can we create a facility that can offer baseline performance metrics and assess a wide variety of UGV platforms, configurations and weight classes? This paper will present efforts to develop a Systems Integration Laboratory (SIL) environment which will focus on testing and validation of a variety of candidate small and mid-sized unmanned ground vehicles emerging from the various Department of Defense (DOD) UGV programs. A description of the SIL concept for UGVs will be provided as well as how it is being applied in the Intelligent Mobility Robotics program at the US Army Tank-automotive and Armaments Research, Development and Engineering Center (TARDEC). The test areas will be described as well, including the Supervised Navigation Test area, the Modeling and Simulation room and the Hot Bench & Test Integration Room. The paper will also summarize the expected activities planned to be conducted in the lab, the range of unmanned prototype vehicles, which may be tested in this facility, and various modeling and simulation concepts planned for the future. Finally, a description of what makes this research facility unique to DoD will be offered.

INTRODUCTION

The notion of using robotic technology on military ground vehicles has been seriously considered for only during the past 20 years or so. The focus of implementing these unmanned vehicle systems – whether they are on the ground (UGV), in the air (UAV), or underwater (UUV) – is to act as either a force multiplier, or to keep soldiers from performing dangerous missions. These kinds of uses are becoming increasingly important given the increasing social and political demands for a minimum amount of casualties in any military operation. The larger question is how to implement these technologies

efficiently and have the military leadership gain confidence in the ability for using these unmanned vehicles to act as an additional vehicle or weapon for the individual soldier. A major transformation of military tactics and structure will commence once unmanned vehicles become commonplace and are considered a part of the typical military unit.

PRESENT AND FUTURE ROBOTIC CONCEPTS

Unmanned vehicle systems are primarily being pursued in the military services because of their perceived ability to assist the solider in the field in performing assigned tasks in some manner. DOD is developing of an array of unmanned vehicle concepts for three primary reasons. The first reason is to minimize and reduce injury risk to all soldiers. Unmanned systems will increase individual soldier and system survivability by reducing their exposure to battlefield hazards. Mission mine-clearing tasks and information scouts.

The second reason for development is to use these unmanned systems as a force multiplier. Unmanned systems will increase force capabilities in several categories, including lethality, survivability, situational awareness, force protection and sustainability.

Lastly, the biggest potential long-term payoff for development and implementation of UGV systems is the prospect of having a vehicle or weapon that requires only minimal direct human involvement after being assigned a task or mission. This, in fact, is the definition of autonomous mobility – being able to internally be given a task and complete it without human intervention. Unfortunately, this type of system is still to be fully realized [1]. Unmanned systems will likely be very specialized and situational dependent. However, they will have a variation of abilities and will likely be employed across the range of military operations.

Current DOD efforts

Most advanced development projects related to unmanned ground vehicle robotics have been consolidated under the

Joint Robotics Program (JRP) directed by the Office of the Secretary of Defense (OSD). JRP is responsible for launching the development of first generation UGVs and related technologies that are deemed critical to follow-on systems. Their goal is to develop and field a family of unmanned ground vehicle systems for a range of military applications in accordance with user requirements. These vehicles range in size from microbots – which can fit in the palm of a person's hand – to teleoperated tanks, such as the Turretless M-60 Panther, which is currently being used to detonate Anti-Personnel mines in Bosnia.

The long-term goal is to mature critical technologies to move robotic applications forward from teleoperation—where a human directly controls UGV functions — to semi-autonomous and autonomous functioning. Semi-autonomous performance of a UGV is a vehicle that functions with the operator only in a supervisory role and thus able to control multiple UGVs concurrently. Autonomous performance of a UGV would be realized when a set of instructions would be carried out without human interaction during the task.

TARDEC Robotics programs

Demo III is the cornerstone of the Joint Robotics Program UGV Technology Enhancement and Exploitation (UGVTEE), tech-base, program which is managed by the U.S. Army Research Laboratory (ARL) and coordinated through TARDEC. The goal of the Demo III Program is to develop technology required to demonstrate a small, survivable unmanned ground vehicle capable of autonomous operation over rugged terrain as part of a mixed military force containing both manned and unmanned vehicles. The vehicle (see Fig. 1) will weigh approximately 2,500 lbs., and be able to maneuver cross-country at speeds of up to 20 mph during daylight, 10 mph during darkness, and up to 40 mph on roads during daylight. This will allow it to successfully operate with a mixed maneuver force during tactical operations.
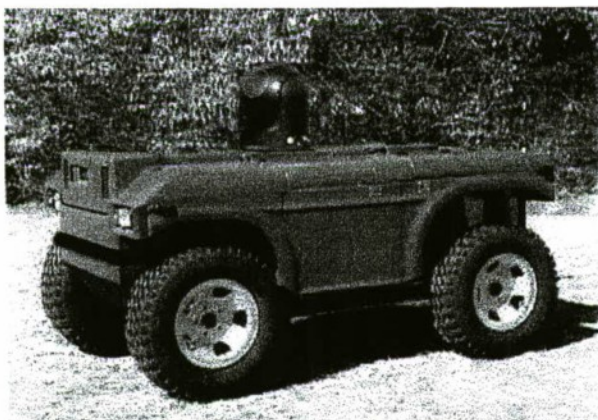
Figure 1 Demo III Experimental Unmanned Vehicle
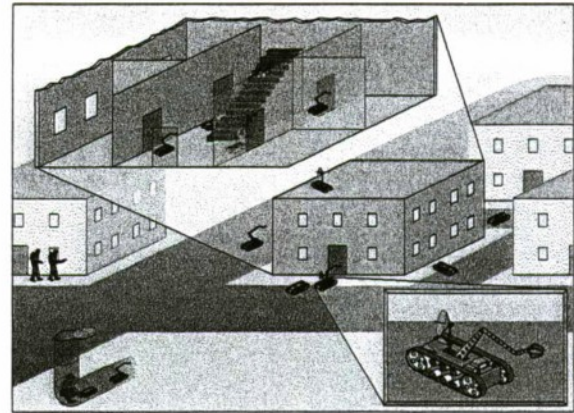
Figure 2 Tactical Mobile Robot MOUT concept:
10 robots to clear 20 rooms in 60 minutes

Tactical Mobile Robotics (TMR) is a DARPA funded program to develop advanced mobile robotic technologies for man portable platforms – targeted at between 10 and 40 pounds -– and integrate them into systems which will be demonstrated in tactical operations. The primary technologies to be developed are robust obstacle negotiation, fault tolerant autonomy, multi-platform perception, and denied area mapping. The program objective is to design, develop, and demonstrate the use of small robot teams in realistic, tactically significant scenarios. Emphasis is placed on man portable platforms working in complex, obstacle intensive environments and denied areas. The primary mission profiles are reconnaissance oriented, with collateral value realized through delivery of specialized mission payloads such as communication relay and mobile obscurant (smoke) projection. A variety of more sensitive mission payloads are also anticipated.

Intelligent Mobility Robotics program

The TARDEC Intelligent Mobility (IM) robotics program is designed to complement and extend survivability and mobility requirements of several DOD robotics programs such as Demo III and Tactical Mobile Robots. Primary program objectives are to develop enhanced mobility and survivability technologies and subsystems for UGV programs in locomotion, machine perception and supervised navigation. IM will also augment present navigational system limitations by developing "smart" drive system technologies via feedback control at its key elements. The work during the first year of the program will focus on vehicles between 100 and 1,500 pound in size. Utah State University has developed two vehicles in the first year of the program. A sub-scale 95-pound omni-directional vehicle (ODV) called T1 (see Fig. 3) was delivered in November 1998, while a mid-scale ODV approximately 1,250 pounds called T2 (see Fig. 4) was demonstrated to the Government in June 1999.
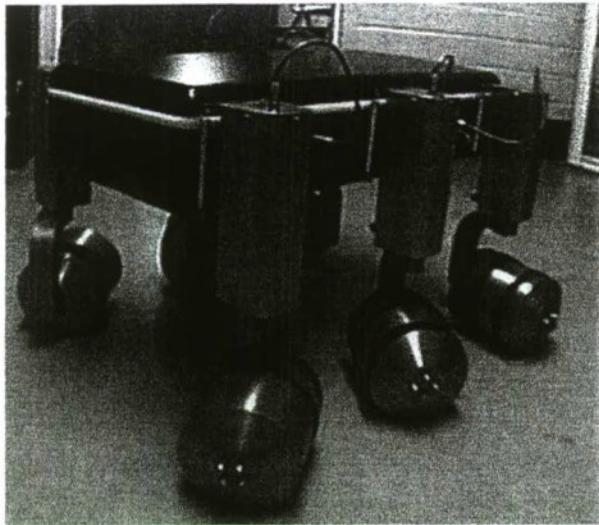
Figure 3 Intelligent Mobility T1 concept vehicle
developed by Utah State University

The essential vehicle elements of both concept vehicles are the running gear configuration, command and control, and path planning. These systems interface with the external world via the soldier interface and digital battlefield communications. The vehicle world model will use a tactical decision aid that exploits situational awareness to determine vehicle operation on the battlefield. Critical elements of the situational awareness algorithms will the synergism between intrinsic mobility and signature management. The tradeoff between moving undetected while avoiding and negotiating obstacles will give tactical flexibility for the UGV to accomplish its mission.

Figure 4 Intelligent Mobility T2 concept vehicle
developed by Utah State University

The sub-scale T1 contains an in-hub electric drive system where each wheel has two degrees of freedom (2-DOF). The wheel hub assembly includes the electric motor, processor board, batteries and an optical slip ring, which connects the drive electronics to the on-board vehicle controller. The aggregation of six independent wheels using an intelligent controller gives the T1 vehicle omni-directional steering (ODS) capability. This unique concept allows the UGV to perform more complex maneuvers such as right angle turns that are currently not possible with an Ackerman steering drive system. Future UGV systems will deploy intelligent wheels with z-axis or 3-DOF allowing the vehicle chassis a 6-DOF motion capability.

The mid-scale T2 also demonstrates the off-road mobility and survivability advantages of cooperative multi-vehicle command and control. A number of iterative changes have been incorporated into the T2 design including moving the batteries and power conditioning equipment out of the wheels and into the vehicle chassis. The T2 is an electric vehicle with no internal combustion engine providing power as in hybrid electric systems. However, the original ODV concept remains intact.

SIL DEVELOPMENT

The concept of a systems integration lab (SIL) has been traditionally used in evaluating complex systems such as vehicle electronics and hydraulics. A SIL is usually set up to allow a "coupling" of prototype system elements. This can be combining two separate hardware technologies, or by joining a piece of hardware with software. For experimentation purposes, the remaining system elements are simulated via software. Traditional SIL applications also include the ability to exercise the system in a mission realistic manner. Prototype hardware is individually brought online until the entire system is functioning together. It is also important that the test vehicle be in proximity to the SIL to allow consistent evaluation of system characteristics during the transition of hardware and software from the SIL into the vehicle.

TARDEC has also designated approximately 4,300 sq.ft. for construction of a SIL to extend the capabilities of the Intelligent Mobility robotics program. This facility will include three major areas for laboratory testing, including a supervised test navigation area, a hot bench and test integration lab, and a modeling and simulation lab.

PROPOSED LAB ACTIVITIES

The SIL was created to perform developmental testing and evaluation of a variety of small-scale UGV systems,

focusing on mobility and survivability technologies. The primary areas for performing these tasks include a supervised navigation test area, a hot bench & test integration room, and a modeling and simulation room. Besides the Intelligent Mobility program, it is foreseen that it will extend the capabilities of a variety of parallel DOD sponsored UGV research efforts, including the aforementioned Demo III and TMR programs.

The laboratory environment will provide ample opportunity to work with these types of vehicles and related technologies in a controlled environment. The Intelligent Mobility program is also investigating opportunities to partner with a military user and/or MOUT (Military Operations in Urban Terrain) facility in order to extend the ability to acquire field data via user interactions with UGVs. Design of the lab has been completed. It is expected that the facility will be operational in FY2001.

Supervised Navigation Test Area

An obstacle course will be constructed in the supervised navigation test area. The course will include several positive and negative obstacles for small-scale unmanned vehicles. Such items as an indoor rock bed, stairs, and a variety of soil types would be incorporated. Several vehicles can then be ran through the course and graded for ability and speed of movement throughout. Testing of navigation and locomotion over these standardized mobility obstacles will be available in this facility. Obstacle detection and avoidance technologies can also be validated here as well.

An endurance test will also be developed. This will gather data on the total distance a UGV travels before losing its ability to move, operating time, and the resulting average speed. Studies may be performed to determine what the nominal operation time might be for any given mission.

Traditional test procedures from legacy (manned) vehicle systems will be used as the starting baseline. Many can be implemented with either little or some adaptation for unmanned ground vehicle systems. From this beginning, test procedures specific to the testing and validation of unmanned vehicle systems may be developed. The goal is to provide data for various types of UGV model validation.

The following are examples of some candidate parameters to be considered for test and validation issues in the test area:

Physical dimensions: Validation and assessment of platform characteristics to evaluate, including but not limited to: length, width, height, ground clearance, weight, and ground pressure.

Tractive Effort: The ability of the vehicle to move its own weight. That is, does the vehicle have the traction and horsepower to transfer power from its powerpack to the ground for mobility purposes? This can be evaluated for both flat terrain and transverse grades/slopes.

Forward/Reverse Speed: The movement of the vehicle and its maximum/minimum speed in forward and reverse. Vehicle acceleration and braking may also be evaluated as well.

Steering: The minimum turning radius of the vehicle and type of steering employed.

Trench crossing: A standardized trench can be developed and built based on the parameters of the vehicles to be tested.

Vertical step: It is foreseen that UGVs involved a MOUT scenario will literally need to climb up and down steps between floors. A standardized step will need to be defined and developed. However, specific vehicle requirements will demand adaptation from that standard.

Fording: Legacy vehicles are required to move through shallow bodies of water. It is not yet known what requirements UGVs will have, but should be considered in the SIL development.

Power management. The location and use of power within the vehicle's operational budget is essential to the completion of the planned missions. The SIL may provide an opportunity to evaluate power generation, power storage, and power distribution.

Actuators/Resolvers: Mechanical devices to move and devices to sense movement of those devices need to be tested in the SIL, particularly electrical and mechanical interfaces.

Suspension: Vibration effects of the suspension will degrade sensor performance. Therefore, it will be necessary to evaluate its performance.

Route following: A variety of sensor inputs such as navigation, stereo vision, map database should be evaluated for performance.

Communications: System internal and external communications are essential to UGV functions: teleoperation and data bus for example. The SIL must be given the facility equipment and structure to allow for the

safe operation of external communications equipment during UGV performance evaluation.

These are only a cursory list developed of the types of parameters to test and evaluate within the SIL [5]. The navigation test area will also be used to evaluate sensor suites for their perception performance on nominal UGVs. In this application, the sensors will be judged on their abilities to assist the vehicles in obstacle avoidance and obstacle detection.

Hot Bench Test & Integration Room

The Hot Bench and Test Integration room is envisioned as a place where researchers investigate individual components and subsystems of small-scale UGVs. An individual piece of vehicle hardware may be highlighted within a virtual UGV system, for example. Once isolated, it can be tested for performance off of the vehicle. Furthermore, new and unique technologies may be integrated in a similar fashion and its performance compared before it may be added to the UGV itself. This offers the ability to evaluate a variety of unique prototype technologies for increasing the mobility and survivability of small-scale UGVs.

The area is also being considered to be a tool for both government and contractors to join together and test the compatabilities of variously independent subsystems. Since many contractors are pursuing the same goal, it is important for the Government to provide a facility to merge hardware and software concepts.

Modeling and Simulation Room

There are three basic areas of work envisioned for this aspect of the SIL. First, is the development of new engineering models. An example of this might be developing UGV-based terrain models for both small-scale and mid-scale prototypes and applying the results to assist the user community in determining how to strategically incorporate UGVs into their missions.

Intelligent Mobility robotics personnel have already begun to review several mobility models such as NRMM to provide additional data. Most of the existing data has been developed for legacy vehicles. However, validation tasks have not been performed for vehicles roughly 1,000 pounds or less. They also have not considered anything but man-in-the-loop scenarios. New models are anticipated to develop as a result from this effort.

Secondly, is the chance to validate control algorithms that will come along from developers. We anticipate using the Modeling and Simulation area in conjunction with the Supervised Navigation Test area to perform these tasks.

Lastly is the potential for wargaming simulations. Since UGVs are a new vehicle class, in a sense, little is known in the user community to how to apply these best in accomplishing a variety of mission tasks.

To assist the user, the Intelligent Mobility robotics program is actively seeking collaboration with the US Army Infantry School at Ft. Benning, GA. Ft. Benning offers the McKenna MOUT site for conducting infantry exercises. McKenna has seven buildings wired for video, sound and continual monitoring. Furthermore, it has an adjoining computer lab facility in place to track movement of soldier activity for experiment and review.

Discussion is underway to have UGVs which pass muster through the Supervised Navigation test area operate in the MOUT facility and transmit the resultant data back to Intelligent Mobility for analysis. A potential mission would be to use the field data to assess how and if infantry goals for incorporation of UGVs are reachable.

POTENTIAL COLLABORATIONS

There are several potential collaborations for the Intelligent Mobility SIL. This includes the previously mentioned the US Army Infantry School at Ft. Benning, GA and assisting the user in updating the ORD for Man Portable Robotics Systems, which are classified in the 10-40 pound platform range.

TMR offers the best opportunity for gathering a handful of different platforms in the near future. TARDEC has been named as the TMR robot repository, and as such, needs to provide storage and maintenance of the back-packable UGVs. The Intelligent Mobility SIL offers a place not only for storage, but ample opportunity to perform additional test and validation on the systems.

As part of the Demo III Concerted Technology Thrust, the SIL may provide unique technologies such as the Utah State omni-directional steering or the Intelligent Wheel system. There may be opportunity to merge some systems within the Hot Bench and Test Integration room.

Other potential collaborations include dual-use applications sponsored by the National Automotive Center (NAC) for the 21st Century Truck; UGV/S JPO efforts, and the potential for miscellaneous contractor work and Cooperative Research and Development Agreements (CRDAs).

REFERENCES

[1] "Unmanned Systems Concepts in the Battlespace" briefing to TRADOC Robotics Conference, 7 Feb 1998.

[2] "US Infantry Mission Brief," COL Frank Stone, Infantry Director, Combat Developments, Ft. Benning, GA, April 1999.

[3] "Concept of Employment for Unmanned Systems," Training and Doctrine Command (TRADOC) Whitepaper, Ft. Monroe, VA., 25 May 1999.

[4] "Concept for Robotic Operations," US Army Engineer School Whitepaper, Ft. Leonard Wood, MO, 14 April 1997.

[5] Robotics Systems Integration Laboratory White Paper, Steve Czerniak, General Dynamics Land Systems, Sterling Heights, MI, 1999.

[6] FY1999 UGV Master Plan, Joint Robotics Program, Office of Secretary of Defense (OSD), http://www.jointrobotics.com.

[7] G.R. Gerhart, R.C. Goetz, D. Gorsich, "Intelligent mobility for robotic vehicles in the army after next," *SPIE Proceedings 1999,* Vol. 3693-15.

APPLICATION OF TIRESOFT FOR THERMAL IMAGE INSPECTION OF DEFECTIVE MILITARY TIRES

Douglas R. Miller,
Project Manager, Thermal Imaging Inspection Station (TIIS)
U.S. Army Tank-automotive and Armaments Command (TACOM)
National Automotive Center
Warren, Michigan

Ferdinand H. Zegel
Radian, Incorporated
Alexandria, Virginia

William R. Reynolds
Signature Research, Inc.
Calumet, Michigan

Randy J. Houle
Signature Research, Inc.
Calumet, Michigan

INTRODUCTION

Two military vehicles assigned to the Arkansas National Guard had accidents, resulting in one fatality, when a tire failed on each vehicle. Representatives of the U.S. Army TACOM Tank Automotive Research, Development and Engineering Center (TARDEC), in conjunction with Radian, Incorporated, arranged for the evaluation of additional tires from the accident vehicles. Testing was conducted at the TARDEC Tire Test Research Laboratory, utilizing thermal imaging technology/analysis developed under a dual military/commercial use project sponsored by the U.S. Army National Automotive Center, located in Warren, MI.

This paper will review the concept, technology approach and application results of the National Automotive Center sponsored dual use initiative, identified as the Thermal Imaging Inspection Station (TIIS), in conducting a real-time assessment of the condition and structural integrity of military truck tires. Particular emphasis will be placed on the analytical software developed by Signature Research, Incorporated in support of the project. This program, identified as "TireSoft", is the basis for the accurate, real time analytical capability of the TIIS system. The application approach and the mathematical basis for the program will be discussed in detail in the document, along with the specific evaluations and results achieved to date.

BACKGROUND

The National Automotive Center(NAC) is funding a Dual-Use Application Program (DUAP) to design and develop a Thermal- Imaging Inspection Station (TIIS) for the preventive detection of flaws or impending failure of bearings, brakes, and tires of military and commercial vehicles with Radian Inc and Signature Research as the performing contractors. These tests were performed on the large tire dynamometers located at Team Truck facilities at the Tank Command in Warren ,Mi. Aware that the NAC-Radian had been testing at this facility deterring tire signatures for the HMMWV tire, Team Truck requested that with the corporation of NAC , Radian,inc use the model thermal-imaging equipment, the improved software and their thermal physics expertise to participate in these tests to determine if one can resolve the cause of the failures might be due to some defect in the construction engineering or manufacture of the tire. Signature Research was asked to modify the original TIRESOFT program in order to make it more calibratable and less influenced by the instrument's gain and contrast controls so that any statistics developed by the software can be carried over into different experiments without loss of confidence in the results. This paper will describe the process we used to develop an in-situ calibration scheme whereby in-scene calibration targets can be used with confidence. Then we

will describe some of the results we obtained during the course of testing a variety of tires both good and bad, new and used. And the insight that came to us as we reviewed the statistical data that was gathered.

DEVELOPMENTAL WORK ON THE SOFTWARE: "TireSoft"

TireSoft is a semi-autonomous application for detecting and analyzing the temperature of tires, brake areas and bearing areas as they move through the field of view (FOV) of the imaging sensor. TireSoft is integrated with the thermal sensor package creating a system for detecting tire defects.

In practice the system is placed so that the tire sidewalls of moving vehicles will pass by the FOV. After minimal user set-up the TireSoft application continually analyzes the flow of images. A rolling wheel passing through the FOV will be autonomously detected and analyzed. The sidewall, brake area and bearing areas are all segmented in the software. If a fault is detected in any one of the three areas the image will be stored and an alarm will be set.

TireSoft was jointly developed by Signature Research, Inc. and Radian. TireSoft is a C++ MFC application.

SEGMENTATION OF THE IMAGE INTO REGIONS

Initially the image is thresholded, meaning pixels with values between two tolerances are set to unity, and the others set to zero. Those set to unity are said to be "tagged", and these are to be partitioned into connected sets which are to be processed. Ultimately each connected region will have its own tag value.

The routine assumes that the image is being processed one line at a time (left to right) from top to bottom. Basically, it looks at three tagged pixels: those to the West, NorthWest, and North. It assigns to the current pixel the value of the minimum
tag of the three tagged pixels W, NW, & W. However, the MinTag functions will look for the minimum tag connected to N, W, or NW; not just the minimum of N, W, and NW. Thus for example suppose W = 400, N = 300, & NW = 350. The current pixel will not necessarily be tagged with 300 because 400 may be connected to a tag value like 200. The number of tags is the number of · connected sets. However, these sets are unioned (by the function "FindConnectedRegions") if connected, or omitted (by the function "EliminateUnusableRegions") if failing further tests. The objective is to obtain the smallest number of sets which are connected, disjoint from other connected sets and viable.

The function "FindConnectedRegions" scans through the Tag Buffer and, for each nonzero tag, replaces it with the associated low tag. It also builds a list of connected areas. The resulting connected regions are bracketed by bounding rectangles. The height and width of the bounding rectangle are then compared to known tolerances. If these dimensions deviate too dramatically from the tolerances, the region is deleted and no analysis is performed on it. This is executed in function "EliminateUnusableRegions".

SEARCHING REGIONS FOR TIRES

Once the set of viable regions is defined, the function "LocateTire" iterates through the list of regions and looks for tire-like features. This means circular outer boundaries, annular sectors and concentric circles (for the tire, wheel and hub). The mechanism for accomplishing this search for geometric characteristics is to start on the region's bounding rectangle. On this rectangle thirteen points are computed: the one third, one half and two third distances along the left, right and top sides, as well as the four corners. From each of these points a ray is cast into the rectangle, vertically down from the top, left or right from the sides and at 45 degrees into the rectangle from the corners.

Note that the bottom side of the rectangle is excluded from these calculations. This is because in many instances the bottom of the tires are partially flat, to a degree that can skew the calculations. Each of the rays is tracked until it intersects the region, and as a consequence 13 boundary points for the region are thus obtained.

The first step is to determine if these boundary points, or some subset of them, essentially lie on a circular arc. There are 286 ways to choose 3 points from the 13, and so, since three points on a circular arc determine the arc's center, there exist up to 286 estimates for the hypothetical center. Great care must be taken in analyzing this data since actual tires may often not appear perfectly circular. For example, if the sensor's line of sight to the tire is slightly oblique the tire appears elliptical. Worse yet, another adjacent tire or one across the vehicle may distort the target tire's region.

Calculations are performed to seek out centers that are outliers. Ultimately the best ten are averaged to get the candidate center. The radius is now that of the tire. Circular templates with variable radii are now fit within the outer circle to denote the wheel and hub. This can be done automatically by examining fluctuations in pixel intensity or the user can fit the templates manually.

This laborious procedure was implemented as it performs extremely well in the presence of significant amounts of clutter and noise.

FIRESOFT CALIBRATION

The sensor graylevels are transformed to engineering units using a two point calibration process.

For calibration purposes the emissivity $\in(\delta)$ of all of the scene elements are assumed to be unity. This assumption is based on the premise that the sensor does not have knowledge of the emissivity of each instantaneous field of view, IFOV, and therefore the sources are assumed to be blackbody radiators. The resultant measured temperature is an *apparent* temperature.

The general equation for transforming gray levels, η to radiance values, $L(\eta)$, is a linear relationship where

$$L(\eta) = m\eta + b$$

over the spectral bandpass, $R(\delta)$, of the sensor. That is,

$$L(\eta) = \int_{0}^{4} R(\delta)L(\delta,T)d\delta \qquad (1)$$

for a zero distance source and where $L(\delta,T)$ is the Planck function. The total spectral emission by a blackbody per unit area *into the hemisphere* is given by the Planck function

$$L_{\delta b}(\delta,T) = B\alpha_{\delta b}(\delta,T) = \frac{2BC_1}{\delta^5(B^{C2/\delta T} - 1)} \qquad (2)$$

where $C_1 = 0.59544 \times 10^8 \, W \cdot \Phi m^4/m^2$ and $C_2 = 14,388 \Phi m \cdot {}^\circ K$

For two black bodies imaged on the same settings at sensor gain setting G_1 and sensor offset setting V_1 and the atmospheric transmission is unity (short range approximation) then over the bandpass of the sensor

$$L_{b1}(\eta) = m_V \eta_1 + b_V$$

and

$$L_{b2}(\eta) = m_V \eta_2 + b_V$$
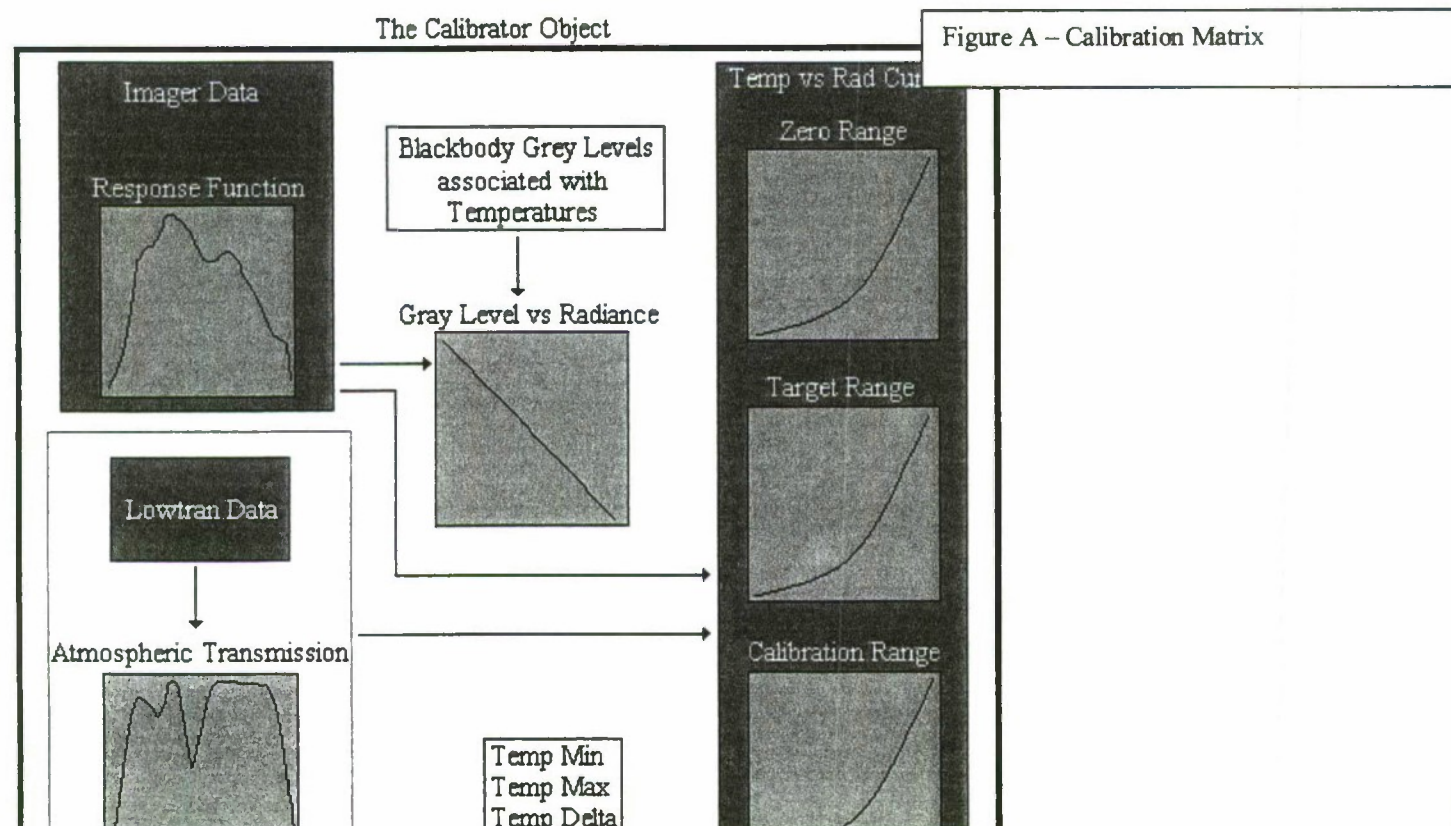
where values for $L_{b1}(\eta)$ and $L_{b2}(\eta)$ are computed using equation (1). Solving for m_V and b_V

$$m_V = \frac{L_{b2} - L_{b1}}{\eta_2 - \eta_1}, \qquad b_V = \frac{\eta_2 L_{b1} - \eta_1 L_{b2}}{\eta_2 - \eta_1} \qquad (3)$$

In summary the calibration procedure shown in the right half of figure A is as follows.

The Calibrator Object

Figure A – Calibration Matrix

TEST METHODOLOGY

TEST TIRES

The established thermal signature of a new, high-quality HMMWV tire determined by the NAC-
Radian TIIS program was designated as a "control", so that results of tests with the 5-ton truck tires could be compared with those for the HMMWV tire. The HMMWV tire is similar in construction, i.e. it is a radial tire, but differs from the truck tires in size.

Two 5-ton wrecker truck tires were tested: one was from the right rear and the other from the right intermediate rear of the M939A2 that had a tire failure in June 1998. To the naked eye, both test tires appeared to be in excellent condition: treads were not overworn and no cuts or abrasions were visible. They had a recorded mileage of 7,890 miles of usage. Both were Michelin X-radial tires with the characteristic heavy, military, off-road bar-type tread. The control and test tires are identified in Table 1.

TIRE DYNAMOMETER

The tire dynamometer (Figure 1) is located in a testing chamber and is operated remotely from the adjacent control room during testing. Steel-wired windows between the control room and the test chamber permit operators and observers to monitor the tests. A video camera mounted in the test chamber permits remote, real-time observation via video monitor and it also preserves a record of each test on videotape.

Tire Test Research Lab's tire-testing apparatus comprises three major components: a mounted steel drum, a tire mount, and a hydraulic ram.

Ambient temperatures in the test chamber were:

Control Tire	Start: 79°F	End: 81°F
Tire 1	Start: 72°F	End: 75°F
Tire 2	Start: 77°F	End: 83°F

STEEL DRUM

A large, steel drum, or flywheel, with a flat, unpolished steel surface is mounted so that it may be
rotated by a heavy-duty electric motor (mounted on the framework above). Rotation speed is adjusted by the gears of the motor. The drum, mounting frame, and motor can be seen in the center of Figure 1.

Tire Mount

The test tire is fitted with a special hub for the axle mount that holds the tire in place between the rotating drum and the hydraulic ram. The hub is bolted to the axle. The axle is mounted in an arm that extends from, and is part of the hydraulic ram.

Hydraulic Ram

Tire loading is accomplished by a hydraulic ram that forces the tire against the drum as it is rotating. The ram is calibrated in total pounds of pressure, or load.

The ram is a large framework, with a guard plate mounted on top, that moves forward as one unit. The ram is shown here from the point of view of the "free" side of the tire to be tested. It is the arm of the ram that forces the tire against the drum.

THERMAL IMAGING INSPECTION STATION (TIIS)

The TIIS as currently configured comprises a thermal-imaging system, a calibration apparatus, a computer system, a video monitor and recorder, and a specially developed software package.

Thermal-Imaging System

The thermal-image acquisition device used for this test was a Cincinnati Electronics IRRIS-160ST Indium-Antimonide infrared focal-plane-array camera adapted for the TIIS. It was equipped with a 50-mm KRS-5 lens. The other component of the thermal imaging system was a remote control keypad linked to the camera via cable. The camera was placed on a tripod, in the test chamber, and located at the closest distance where the entire side of the rotating tire could be captured in the image. During tests on the tire dynamometer, however, a portion of the tire is blocked by the framework of the ram. This limitation is consistent for all tests, including those of the control tire. Also, the rotational velocity of the tire during the 33 milliseconds it takes to capture an image does bias the data. The net effect is to slightly smear the image pixels during any exposure. The manner of coping with this limitation is a data analysis issue and is discussed below.

The camera is capable of capturing 30 images/second; however, the computer and software can only record 12 images/minute. (Note: A faster computer could capture more images/minute.) An image consists of 19,200 pixels. Each pixel has an intensity value between 0 and 4,096. These values are translated into temperature readings by software within the camera.

Blackbody Calibration Source

A Mikron 310 Blackbody Calibration Source was placed within the camera image frame, in the area of the ram that was blocking the view of the tire. The Blackbody—a small box that houses the controllers, heater, and cooler—is part of the thermal scene. The Blackbody was set to provide a specific, consistent source of thermal heat,

132

accurate to ∀ 0.2° C after stabilization. (*Note: it takes approximately 30 minutes to stabilize the Blackbody.*) The operator of the TIIS used it to calibrate the system by means of computer.

Computer System

The computer was a Micro Systems 133MHZ Pentium 1 PC with a 1.2 Gb hard drive and a MicroScan 5V/AOI monitor. The system included an Iomega 100 external Zip drive connected between the parallel port and the Hewlett-Packard Desk Jet 820Cse color printer.

Images were captured remotely by the computer system, which was located outside the test, outside one of the garage-type doors located behind the hydraulic-ram end of the several tire dynamometers in the chamber (on the opposite end of the test chamber from the Tire Test Research Lab control room). TIIS operators had no visual access to the chamber during testing.

The computer monitor displays the thermal images in real-time, at a 30-Hertz video rate, while they are processed in the computer.

Video Monitor and Recorder

In addition to the computer system, a JVC S-5300U S-VHS tape recorder and a Magnavox color monitor were connected to the system for real-time monitoring and recording. Playback from the tape recorder degrades the image quality approximately 50% and is used simply as a visual record of the tests.

Software

The software of the TIIS is a Radian Inc./Signature Research, Inc. jointly developed tire diagnostic software called "TireSoft". TireSoft has the unique ability to perform a statistical analysis on an exact image area of the tire being tested. The software is designed to divide an image into three donut-shaped concentric areas with increasing radii. These areas are designated Zone 1 (the hub), Zone 2 (the wheel), and Zone 3 (the tire).

The boundaries are user defined; they can be sized to fit the tire exactly. Unwanted elements in a particular zone can be excised and calculations made from the balance of the pixels in the zone. TireSoft performs four functions on each zone. The software:

- reads the one pixel that has the highest temperature,
- reads the one pixel that has the lowest temperature,
- calculates the mean temperature of all pixels within the zone, and
- calculates the standard deviation.

All data reported for the tire tests was recorded from Zone 3.

TIRE TESTS

The recommended inflation pressure (psi) and load weight (lbs) are imprinted on the sidewall of every tire (Table 1 details the physical characteristics of the Control and Test Tires). The control tire had been tested at the recommended pressure and a load of 2,000 lbs. Tire 1 was tested at the recommended pressure and load. Table 2 is a summary of the values for the variables in the tests described below.

Test of Control Tire

The control tire had been tested at a road speed of 35 mph with a load of 2,000 lbs and an inflation pressure of 50 psi. The thermal-image data for this control tire was gathered after 1 hour of running time. An image was captured in digital form every 60 seconds and stored directly on the portable Zip drive.

Test of Tire 1

Tire 1 was tested at a road speed of 35mph with a load of 9,920 lbs and an inflation pressure of 100 psi. It was tested for 1 hour and 27 minutes before it failed. The thermal-image data for Tire 1 was gathered every 60 seconds.

Test of Tire 2

Tire 2 was initially run at a road speed of 35 mph with a load of 9,350 lbs (rated load is 9,920 lbs) and an inflation pressure of 80 psi (standard for this tire is 100 psi). It was tested for 58 minutes at that setting. At the 58-minute mark, the road speed was increased to 45 mph. After a further 29 minutes (at 1 hour and 27 minutes), the road speed was increased to 55 mph. After an additional 30 minutes (at 1 hour and 53 minutes), the load was increased to 9,900 lbs. The tire failed 9 minutes later (at 2 hours and 2 minutes). Thermal-image data for Tire 2 was also gathered every 60 seconds.

STATISTICAL ANALYSES

In order to account for the 33-millisecond time lag to record an image of a rotating tire, Radian integrates data from a number of frames and then analyzes the pixel variance in the images. Also, from experience gained during the ongoing NAC-Radian program, it had been learned that a series of 30 images captured at a rate of one per minute would record many instances of all areas of the tire being testing, including those blocked out by the dynamometer ram. The mean standard deviation (SD) was calculated for 30-image data samples on all three

tires. Then, the mean SD of the control tire was compared with the mean SD of each of the test tires.

RESULTS

The term "signature" is used in thermal imagery to connote an identifiable thermal pattern. Most often the term has been applied to the identification of whole objects, such as various types of tanks and trucks, or to distinguish the presence of people and animals in dark environments. The term is used in describing thermal images from tire tests, not to identify tire types, but to describe the visual pattern that becomes apparent as the tire heats up during testing.

A good tire has a thermal pattern, or signature, that shows that the temperature is reasonably consistent after the tire reaches its operating temperature. By "operating temperature", we mean that the tire has reached a temperature that remains stable over 3–4 minutes of testing. During the ongoing NAC-Radian program, in a series of tests of new tires of various types—including several like HMMWV radial tires, one of which was the control tire—it was established that good tires generally stabilize after about 30 minutes of testing.

In that same series of tests, problem tires exhibited a very different pattern. For example, a tire that had broken cords, not apparent externally, showed a thermal pattern that diagnosed the problem (Figure 2). It should be noted at this time that the criteria for selection of tires included the following: tires must be operational, they must hold air, and they had to be driveable, in other words, tire tests were not run on tires with any readily apparent, visual or functional flaws that would render them incapable of operation. Tires that are damaged, flawed, or unevenly worn may show "hot spots" or "cold spots" that cause the temperature to fluctuate or continue to rise rather than to stabilize, resulting in an elevated standard deviation. Another test of a tire with a bulge, visible only on close examination, might have been expected to reveal a hot spot where the bulge appeared on the series of images, but instead the bulge was indicated by a "cold spot" in the signature (Figure 3). Preliminary analyses of these tires shows a mean SD>1, comparable with the degree of instability of the tire.

THERMAL-IMAGE DATA FROM TIRE TESTS

Data showing the mean pixel value, standard deviation, and maximum and minimal pixel values are available for each image recorded. The Blackbody was calibrated to 40°C before the beginning of each test.

Test of Control Tire

Baseline data on the control tire, a new HMMWV radial tire shown to be of good quality, had been obtained during earlier tests at the Team Truck facility in October 1998. Thermal measurements revealed a steady 37°± 1°C as the mean operating temperature of the control tire.

Figure 4 is a typical image from the test of the control tire, captured after the sidewall temperature had stabilized. It illustrates the signature of the control tire after 60 minutes of testing under the conditions described in the *Methods* section above. The smaller circle in the lower right-hand corner shows the thermal image of 40°C. There are no discernible hot or cold spots. That the color is distributed rather evenly over the entire image indicates that there is little fluctuation in the temperatures within the area.

Test of Tire 1

Tire 1 revealed "hot spots" from early on in the tests, as soon as the tire began to warm up. By 09:58, 5 minutes after the start of the test (Figure 5), the tire showed evidence of a signature not found in a good tire. A band of higher temperature is noticeable in the sidewall. By 10:03, 9 minutes after the start (Figure 5), hot spots are also evident in the tread. At 10:30, the hot spots are more pronounced. The explosion that occurred when the tire failed, at 10:42, was viewed on the real-time monitors. Rupture is first visible to the camera in the image taken at 10:43 (Figure 6).

Test of Tire 2

The thermal image of Tire 2 also showed hot spots after only 5 minutes of testing. Figure 7 shows a wide band of heat in the sidewall as well as heat spots in the tread. After 1 hour and 51 minutes, the sidewall can be seen to be excessively hot.

STATISTICAL RESULTS

The mean temperature of the control tire warmed <1°C during that period, while both test tires heated more rapidly. Also note the temperature variability, as expressed in SD, increased dramatically for the test tires (Chart 1).

Table 3 is a 30-image sample of data on the control tire, captured after 82 minutes of testing (11:01–11:30). The timing of the 30-frame sample was chosen to illustrate the thermal signature of a good tire well after the tire had warmed to a stable operating temperature. The mean SD was 0.552. The SD was not >1 during any of the 30 minutes of testing.

Table 4 is the thermal-image data accumulated during the final 30 minutes before Tire 1 failed. The 7-sec time gap between 10:19:37 and 10:20:44 shows an interruption in

data capture while the operator was reviewing the data. Again, at the first indication of possible tire failure, the operator reprogrammed the data capture rate from 1 frame/minute to 1 frame/5 sec, which caused the time gap between 10:40:44 and 10:44:11. (*Note that the consistent data capture for the control tire is for a period well after warm up, during which time the operator had had ample time to adjust contrast and other variables to assure the most revealing image for that tire. And, when the temperature is relatively stable and no problem area is presenting, the operator does not have a need to examine the data or to adjust the image in an effort to "see" and record more clearly what may be happening.*) These operator actions had no other effect on the data than "time data was stored".

The mean SD was 2.83. Data for two images following the failure are given to show the effects of the failure.

Table 5 shows the thermal-image data accumulated during the final 30 minutes before Tire 2 failed. Variations in the number of seconds from one image frame to the next were caused by interruptions in data capture while the operator was reviewing data.

The mean SD at failure was 4.38. Data for two images following the failure are given to show the effects of the failure.

Figure 8 is a graphic representation of the variations in standard deviation for each thermal image in the sample data of the control and test tires.

CONCLUSIONS AND RECOMMENDATIONS

CONCLUSIONS AND LONG TERM PROGRAM VISIONS

There is a symmetry between the tire image and the analysis of thermal distribution as expressed in standard deviation, i.e., the thermal image that displays elevated

heat zones visible on the TIIS monitor will have a mean SD>1 while images displaying an even heat distribution will have a mean SD<1. From our limited tests, we conclude that a mean SD>1 is associated with a faulty tire.

PROGRAM VISION

In cooperation with industry the potential exists to develop and build a low cost imaging system with adequate memory chip capacity to accommodate the enhanced "Tire-Soft" program. This tool would be placed in the hands of the military or commercial maintenance supervisor or maintenance/inspection personnel, and be used by these individuals as a tool to conduct efficient, cost effective, real-time maintenance analysis of truck and trailer components. This approach has significant implications with respect to military or commercial fleet management/operating costs, for vehicle readiness and reduction of downtime, and probably most significantly, for vehicle operator and roadway safety. This packaging of commercially available technology, in combination with Radian, Signature Research, Inc., and the University of Tennessee developing custom software, designed for ease of use and to be user friendly and give an accurate real-time interpretation of the captured component thermal image, will take advantage of technology development by identifying and precluding truck and trailer component anomalies prior to catastrophic failure. The National Automotive Center, along with Radian, Inc., and our industry and academia partners will continue development of the Thermal Imaging Inspection Station to incorporate the latest technology and cost advantages in maturing this system for both military and commercial marketplaces.

REFERENCES

Miller et al., "Thermal Image Inspection of Tires for a 5-ton Military Wrecker",
Final Report: December, 1998 NAC/Radian Report

Table of Illustrations

List of Tables

Figure 1 - Tire Dynamometer Test Cell
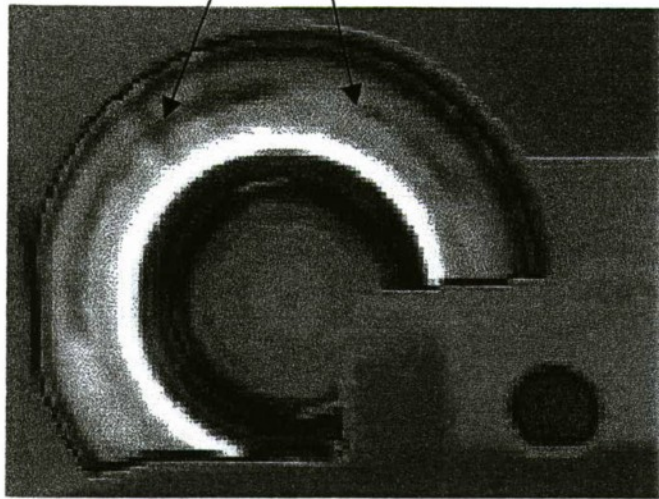
Figure 2. Thermal image of a tire with broken cords, not apparent externally. SD is 1.5 for the temperature range in this image.
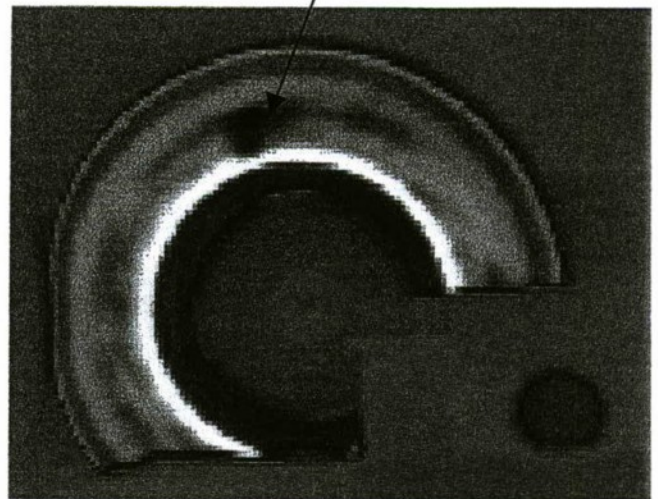
Figure 3. Thermal image of a tire with a bulge in the sidewall, visible only on close examination. SD is 1.3 for the temperature range in this image.

Figure 4. Typical thermal image of a good radial tire at operating temperature, the control tire.
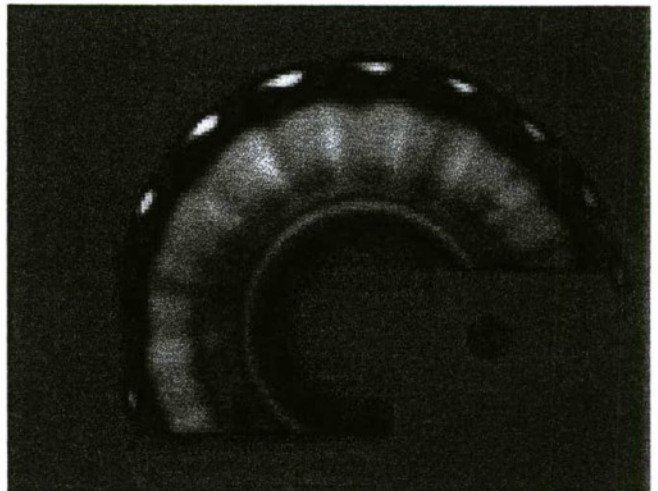
Figure 5. Thermal image of Test Tire 1 after 37 minutes of testing at rated load and psi.
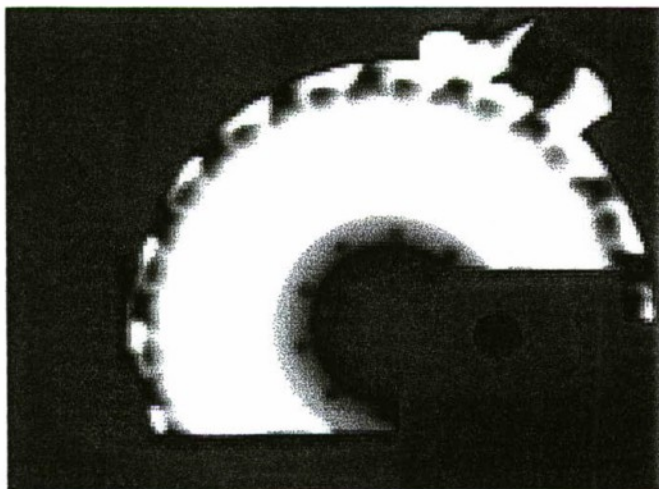
Figure 6. Thermal image of Test Tire 1 after tire failed. 138
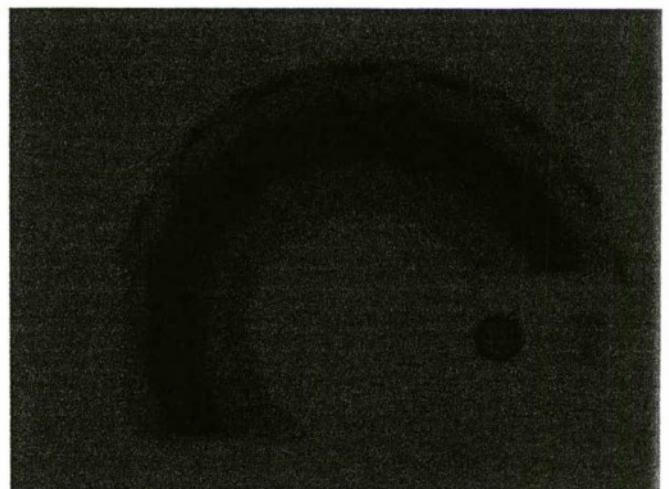
Figure 7. Thermal image of Test Tire 2 after
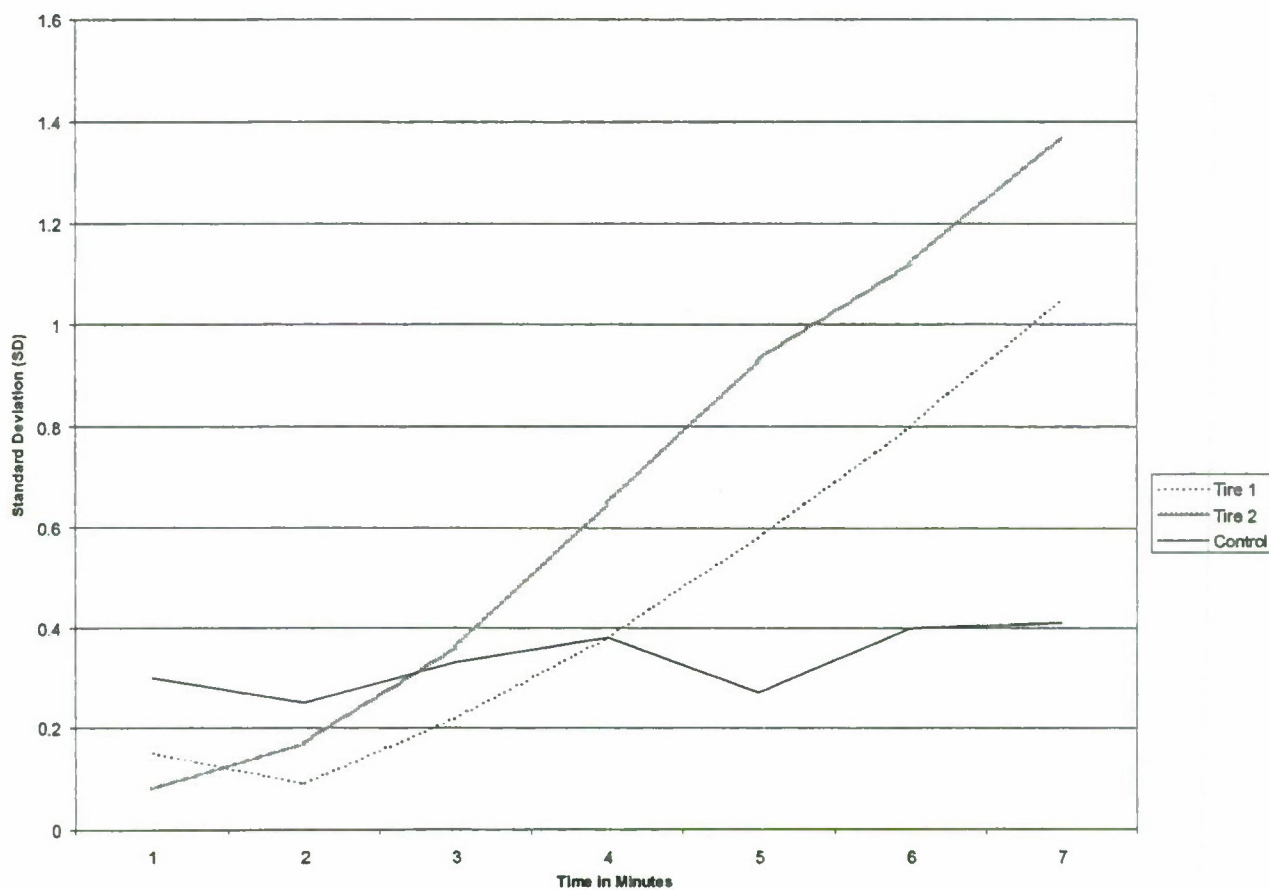
Figure 8. Chart of the standard deviations of thermal image data and test tires

<table>
<tr><td colspan="5" align="center">TABLE 1
Characteristics of Control and Test Tires</td></tr>
<tr><td>Tire</td><td>Type
DOT#</td><td>Manufacturer</td><td>Size</td><td>Max Load
Single</td></tr>
<tr><td>Control</td><td>Radial
59F-4NEV087</td><td>Goodyear</td><td>37x12.50R16.5LT</td><td>3,850lbs@50psi</td></tr>
<tr><td>Tire 1</td><td>X-Radial
VVI-85325A*
Load Range J
Tubeless</td><td>Michelin</td><td>14.00R20</td><td>9,920lbs@100psi</td></tr>
<tr><td>Tire 2</td><td>X-Radial
VVI-85203A*
Load Range J
Tubeless</td><td>Michelin</td><td>14.00R20</td><td>9,920lbs@100psi</td></tr>
<tr><td colspan="5">*Mfg Number – DOT # same for both tires – HC2w1tbx400</td></tr>
</table>

TABLE 2
Summary of Tire Test Variables

Tire	Road Speed (rpm)	Load Weight (lbs)	Inflation Pressure (psi)	Duration (min)	Thermal-Image Rate #/sec
Control	35	2,000	50	116	1/60
Tire 1	35	9,920	100	87_failure_	1/60
Tire 2	35	9,350	80	58	1/60
	45	9,350	80	29(*T=87)	1/60
	55	9,350	80	30(T=113)	1/60
	55	9,900	80	9(T=122)_failure_	1/60

*T = total time from start of test

TABLE 3
A Thirty Thermal-image Sample of Control Tire Data, Captured after 82 Minutes of Testing (11:01-11:30)

Temperature °C

Date, Time	T*	Min.	Max.	Mean	SD
10/22/98 11:01:01	82	34.51	41.78	37.63	0.54
10/22/98 11:02:01		35.00	40.88	37.67	0.51
10/22/98 11:03:01		34.55	41.42	37.61	0.57
10/22/98 11:04:01		34.60	41.75	37.60	0.51
10/22/98 11:05:01		34.48	41.64	37.67	0.55
10/22/98 11:06:01	87	35.27	41.94	37.70	0.55
10/22/98 11:07:01		34.52	41.96	37.63	0.54
10/22/98 11:08:01		34.58	41.90	37.72	0.55
10/22/98 11:09:01		35.16	41.31	37.63	0.57
10/22/98 11:10:01		35.28	41.72	37.70	0.55
10/22/98 11:11:01	92	34.62	41.84	37.74	0.55
10/22/98 11:12:01		34.63	41.53	37.79	0.57
10/22/98 11:13:01		35.19	41.66	37.77	0.58
10/22/98 11:14:01		34.75	41.75	37.72	0.50
10/22/98 11:15:01		34.56	41.71	37.73	0.59
10/22/98 11:16:01	97	34.55	41.14	37.70	0.57
10/22/98 11:17:01		35.11	41.51	37.75	0.51
10/22/98 11:18:01		34.79	41.51	37.73	0.59
10/22/98 11:19:01		34.99	41.88	37.85	0.53
10/22/98 11:20:01		35.33	42.09	37.79	0.55
10/22/98 11:21:01	102	35.33	42.04	37.85	0.55
10/22/98 11:22:01		34.93	41.96	37.83	0.60
10/22/98 11:23:01		34.66	42.08	37.86	0.55
10/22/98 11:24:01		35.05	40.94	37.89	0.58
10/22/98 11:25:01		35.12	42.05	37.88	0.59
10/22/98 11:26:01	107	34.68	42.25	37.85	0.55
10/22/98 11:27:01		35.15	41.67	37.87	0.50
10/22/98 11:28:01		34.74	41.91	37.93	0.54
10/22/98 11:29:01		34.67	41.21	37.93	0.57
10/22/98 11:30:01	111	35.41	42.13	37.98	0.55

Mean S D 0.552

*T = Time, in minutes, elapsed from start of test

Test Duration: 09:39 to 11:40

TABLE 4
Thermal-image Data Accumulated During the Final 30 Minutes Before Tire 1
Failed, and Data from Two Images after Failure.

Temperature °C

Date, Time	T*	Min.	Max.	Mean	SD
12/16/98 10:13:37	18	35.57	50.26	41.81	2.07
12/16/98 10:14:37		35.90	46.74	42.24	2.28
12/16/98 10:15:37		35.78	46.35	42.24	2.08
12/16/98 10:16:37		35.87	47.42	42.58	2.18
12/16/98 10:17:37		35.53	52.09	42.62	2.26
12/16/98 10:18:37	23	35.32	53.79	43.01	2.42
12/16/98 10:19:37		35.99	47.65	43.09	2.33
12/16/98 10:20:44 **		36.08	48.93	43.35	2.33
12/16/98 10:21:44		36.48	55.33	43.61	2.36
12/16/98 10:22:44		36.31	49.05	44.18	2.48
12/16/98 10:23:44	28	36.66	50.42	44.26	2.54
12/16/98 10:24:44		36.82	57.43	44.28	2.53
12/16/98 10:25:44		35.56	49.15	44.45	3.12
12/16/98 10:26:44		36.69	49.17	44.54	2.66
12/16/98 10:27:44		35.89	49.43	44.69	3.03
12/16/98 10:28:44	33	36.24	58.30	44.97	3.01
12/16/98 10:29:44		36.20	57.04	45.02	3.03
12/16/98 10:30:44		36.11	50.60	45.14	3.17
12/16/98 10:31:44		36.35	52.81	45.18	2.90
12/16/98 10:32:44		36.32	58.38	45.26	3.22
12/16/98 10:33:44	38	36.16	50.79	45.58	3.28
12/16/98 10:34:44		36.71	50.65	45.67	3.16
12/16/98 10:35:44		35.50	59.74	45.74	3.39
12/16/98 10:36:44		35.58	60.33	45.92	3.41
12/16/98 10:37:44		36.06	55.25	46.12	3.57
12/16/98 10:38:44	43	36.36	56.16	46.25	3.20
12/16/98 10:39:44		36.75	61.83	46.70	3.52
12/16/98 10:40:44		37.28	51.80	46.76	3.30
12/16/98 10:42:11 ***		36.24	51.65	46.80	2.60
12/16/98 10:43:36	47	36.39	55.37	47.57	3.35

				Mean	2.826
Tire Fails				SD	
12/16/98		36.64	76.50	48.55	3.72
12/16/98		36.20	77.76	49.89	3.68

*T = Time, in minutes, elapsed from start of test

**Data capture interrupted by TIIS operator to review
first indication of possible tire failure data capture was reprogrammed
from 1 frame/min. to 1 frame/5 sec. This caused the gap between 10:40:44 and 10:42:11

Test Duration: 09:55 to 10:43

TABLE 5
Thermal-image Data Accumulated During the Final 30 Minutes Before Tire 2 Failed, and Data from Two Images after Failure

Temperature °C

Date, Time	T*	Min.	Max.	Mean	SD
12/16/98 16:30:10	32	36.80	55.50	49.70	4.50
12/16/98 16:31:00 **		40.10	57.85	52.29	4.10
12/16/98 16:32:00		39.27	57.94	51.74	4.53
12/16/98 16:33:00		40.10	58.08	52.63	3.81
12/16/98 16:34:00		40.72	57.71	52.56	3.93
12/16/98 16:35:00	37	37.61	58.37	52.74	4.03
12/16/98 16:36:29		40.37	58.46	53.02	3.88
12/16/98 16:37:29		40.26	58.70	53.21	3.97
12/16/98 16:38:29		40.12	58.36	52.56	4.31
12/16/98 16:39:29 **		40.30	58.35	52.79	4.17
12/16/98 16:40:29	42	38.02	59.41	53.14	4.09
12/16/98 16:41:34 **		40.10	58.85	52.87	4.44
12/16/98 16:42:04 **		41.02	58.82	53.19	4.19
12/16/98 16:43:04		40.23	59.35	53.29	4.34
12/16/98 16:44:04		40.10	58.89	53.19	4.27
12/16/98 16:45:04	47	40.66	58.98	53.09	4.47
12/16/98 16:46:04		40.50	58.78	53.31	4.29
12/16/98 16:47:04		40.73	59.17	53.38	4.30
12/16/98 16:48:04		38.65	59.50	52.99	4.59
12/16/98 16:49:04		40.91	59.94	53.64	4.65
12/16/98 16:50:04	52	40.69	60.19	53.90	4.43
**12/16/98 16:51:04		39.22	60.36	54.12	4.35
12/16/98 16:52:04		38.88	60.06	53.68	4.71
12/16/98 16:53:04		38.14	60.79	54.35	4.32
12/16/98 16:54:04		40.44	60.18	54.17	4.48
12/16/98 16:55:04	57	41.00	61.21	54.56	4.87
12/16/98 16:56:44		39.69	60.73	54.14	5.09
12/16/98 16:57:44		40.79	60.85	54.56	4.65
12/16/98 16:58:44		39.67	61.73	54.51	5.23
12/16/98 16:59:44	61	38.46	62.90	55.22	4.58
Tire Fails				Mean	4.38
12/16/98 17:00:14		38.46	78.00	54.34	6.39
12/16/98 17:00:44		43.66	79.30	59.55	5.65

*T = Time, in minutes, elapsed from start of test
**Data capture interrupted by TIIS operator to review data. Load increased from 9350 lbs to 9900 lbs at 16:51

Test Duration: 14:58 to 17:00

142

Bekker's Terramechanics Model for Off-road Vehicle Research

Sean Laughery
Grant Gerhart
Richard Goetz

US Army TARDEC
Warren, MI 48397-5000

ABSTRACT

Bekker's Derived Terramechanics Model (BDTM) is an analytical tool for evaluating vehicle off-road mobility. BDTM has been developed using Bekker's equations for vehicle soil interactions. He developed the bevameter technique to measure mechanical strength characteristics for many soil and snow conditions. This procedure uses seven parameters to describe soil conditions, which differs from the conventional single parameter vehicle cone index methodology used by the NATO Reference Mobility Model (NRMM). NRMM uses the cone penetrometer technique to experimentally measure fine-grained soil mechanical characteristics.

BDTM is in a spreadsheet format, and its primary purpose is to compare mobility characteristics for robotic track and wheeled vehicles under different terrain conditions. Bekker's model is a simple, linear one degree-of- freedom (1-DOF) model, which assumes that in a perfectly cohesive soil (i.e. clay), soil thrust is only a function of contact surface area. The model also assumes that for a perfectly cohesionless or frictional soil (i.e. dry sand), soil thrust is a function of vehicular weight[1]. This paper attempts to compare the mobility characteristics of wheeled vs. track vehicles for different size, weight and terrain conditions.

INTRODUCTION

BDTM was developed as a design tool to compare different types of robotic vehicle mobility performance characteristics. No single vehicular locomotion system has optimal mobility performance under all terrain conditions. Vehicle running gear design always involves design compromises or tradeoffs over a number of mobility factors. Most future Army robotic vehicle platform concepts fall into two broad categories: wheeled and track systems.

Wheeled vehicles are typically more agile and maneuverable than tracked vehicles, but possess higher ground pressures and are therefore less trafficable. Tracked vehicles on the other hand have a lower ground pressure, superior traction and are thus more trafficable. However, they are not as agile or mechanically efficient as their wheeled counterparts due to (typically) larger mass and much larger internal motion resistance.

Both wheeled and tracked vehicles have been successful in negotiating roadways and moderately unstructured off-road terrain. Vehicles with a larger wheelbase, ground clearance and horsepower per weight ratios generally have much better intrinsic mobility performance than smaller systems. A comparison of vehicle types for equal size and weights indicates that wheeled systems are typically superior to track systems in agility, maneuverability, ride quality and terrain damage. Tracked vehicles have distinct advantages relative to stability, ground pressure, maximum vertical slope, and drawbar pull.

Selection of running gear usually becomes a choice between which mobility characteristics are most important for a vehicle's intended mission profile. Ride quality is not as important to unmanned or robotic vehicles unless equipment such as sensors exceed vibration limits or structural loading specifications are exceeded for rough terrain conditions. The vehicle need only have sufficient drawbar pull to transport itself and its payload. Low ground pressure is principally an advantage only in soft soil terrain conditions. Unmanned systems generally weigh less and have a lower ground pressure than the larger manned combat vehicles such as the main battle tanks or infantry fighting vehicles.

Agility and maneuverability are both advantageous for off-road conditions. Ground clearance, maximum slide slope angle and wheelbase are important for difficult obstacle negotiation challenges such as ditch crossings or wide vertical steps. In general a complete systems

analysis is necessary to determine the optimal set of mobility characteristics for a particular mission profile. Unmanned vehicles in particular need a new set of system requirements and represent a separate set of design challenges from their traditionally manned counterparts.

BDTM is thus a modest attempt to examine tradeoffs between different mobility characteristics for wheeled and track vehicles. It is a first-order linear model, which ignores the nonlinear dynamic interactions between the vehicle and its terrain. It does, however, analyze three primary parameters essential to generic mobility: vehicle size, weight and ground pressure. Future systems will vary significantly in these parameters. Since they will also navigate over terrain with large variations in mechanical properties, BDTM provides a useful tool for determining their first-order design characteristics.

THE BEKKER MODEL

The mechanical behavior of soils varies considerably under a wide variety of environmental conditions. For example composition, moisture levels, porosity, temperature, etc., affect bulk soil mechanical behavior relative to vehicle/terrain dynamics. It is also well known that for the same amount mechanical loading, a tracked vehicle may cross soft terrain without considerable slippage, whereas wheels may slip considerably, or simply spin. The amount of slip varies with soil type.

The Bekker model uses the relationship between certain physical soil characteristics and shearing strength to predict vehicle cross-country mobility. Bekker considers wheels and tracks as simple loading surfaces having similar forms, but different lengths and widths. He extrapolates the analogy between soil shear produced by laboratory crawlers to track vehicles as shown in Fig. 1a [1]. When the blocked track is moved relative to the soil mass in the laboratory shear box, the maximum shearing force is not developed instantaneously with the initiation of relative motion. Instead the soil must be compacted to some degree before reaching the final steady state mechanical shearing stress. Thus the track grousers begin slipping before reaching the point of maximum vehicle traction. This transient condition is the basis for Bekker's simple 1-DOF model for vehicle trafficability.
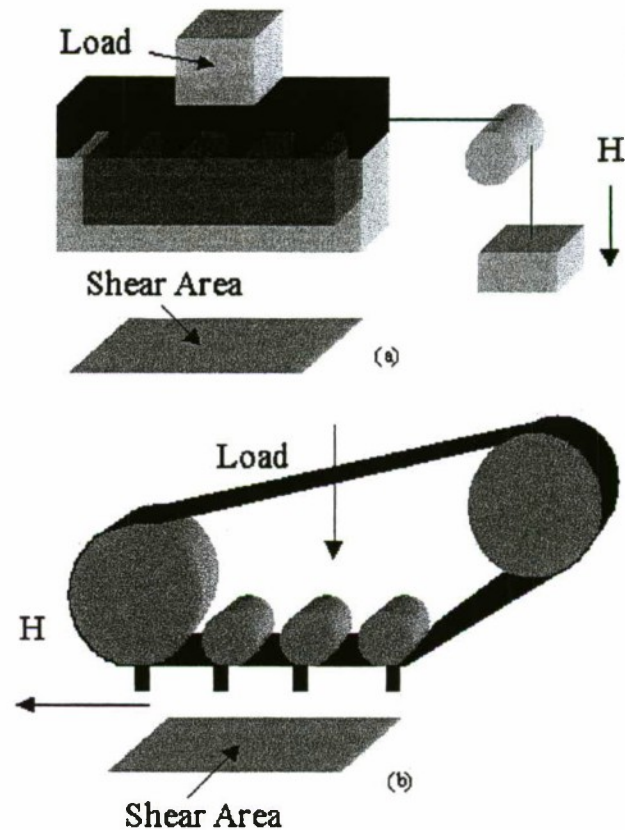

Figure 1   Soil Shear Analogy

The shear stress is the ratio between the vehicle traction force, which is parallel to the soil surface, and the area of the track normal to the surface. This tractive force is opposed by the soil resistance as the grousers slip during the shearing process. The normal loading force of the vehicle compacts the soil, which affects the resistance it exudes against the grousers as the track rotates on the vehicle. In effect the track forces, which push against the soil, generate a soil resistance that is determined by soil type and compaction. Vehicle weight generates ground pressure, which further compacts the soil and alters the soil resistance.

Figure 2b shows a tracked vehicle in motion. A grouser on the track first comes into contact with the ground at position 1. At the moment of first contact no shearing has occurred. As the vehicle moves forward, a shearing force is developed in the lateral direction. The positioning of the grouser begins to slip back pushing the soil and causing a soil distortion (S). As the vehicle continues to move, the amount of soil distortion increases[1].

A. Cohesionless Soil
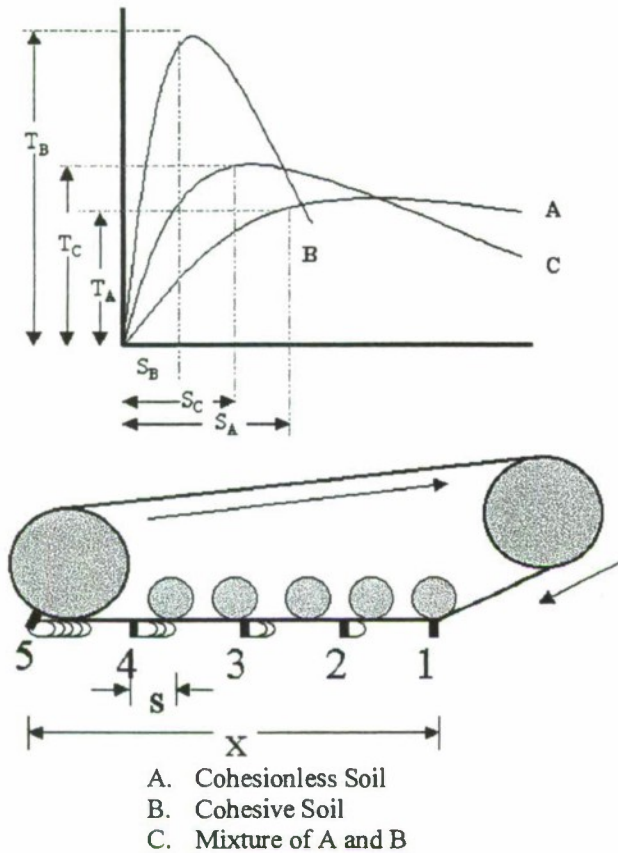B. Cohesive Soil
C. Mixture of A and B

Figure 2   Characteristics of soil deformation

Empirically generated curves in Figure 2a show the motion of soil under shear plotted for three different types of soils[1]. These curves are obtained through empirical data. The curve labeled **A** is for a loose frictional or plastic soil such as wet clay. The shearing strength τ_a of such a soil is reached after the initial period of compaction, which takes place over a distance S_a. After this point the stress remains practically the same irrespective of any slip. Soil **B** consists of a dry coherent mass: dry clay or snow at very low temperatures. This type of soil quickly reaches its maximum shearing strength and then shears off rapidly. The last curve **C** is a soil type that has intermediate properties. Upon reaching a maximum value at a certain slip distance from the origin, it starts to lose its shearing strength but not as rapidly as curve **B**[2].

For modeling purposes, it is critical to come up with a general equation for these curves. The curves in Figure 2a are identical to the displacement (x) and natural time frequency (ωt) of an aperiodic vibration:

$$x = A_1 e^{(-b+\sqrt{(b^2-1)})\omega t} + A_2 e^{(-b-\sqrt{(b^2-1)})\omega t} \quad (1)$$

where b is the coefficient of damping. To write a formula in terms of soil stress (τ) and soil deformation (S),we place $\tau = x$, $K_1 S = \omega t$, and $K_2 = b$ where K_1 and K_2 are coefficients of slippage to get the following result.

$$\tau = A_1 e^{(-K_2+\sqrt{(K_2^2-1)})K_1 S} + A_2 e^{(-K_2-\sqrt{(K_2^2-1)})K_1 S} \quad (2)$$

To determine the coefficients A_1 and A_2 , for slip S = 0 and $\tau = 0$:

$$A_1 + A_2 = 0$$

Also for slip S = 0, $\tau = 0$, and $d\tau/dS = K_3$,

$$d\tau/_{dS} = A_1 e^{(-K_2+\sqrt{(K_2^2-1)})K_1 S}(-K_2 + \sqrt{(K_2^2-1)})K_1 \quad (3)$$
$$+ A_2 e^{(-K_2-\sqrt{(K_2^2-1)})K_1 S}(-K_2 - \sqrt{(K_2^2-1)})K_1$$

$$= K_3$$

A_1 and A_2,

$$A_1 = \frac{K_3}{2K_1\sqrt{(K_2^2-1)}} \quad (4)$$

$$A_2 = -\frac{K_3}{2K_1\sqrt{(K_2^2-1)}} \quad (5)$$

Substituting A_1 and A_2 into equation (2):

$$\tau = \frac{K_3}{2K_1\sqrt{(K_2^2-1)}}(e^{(-K_2+\sqrt{K_2^2-1})K_1 S} - e^{(-K_2-\sqrt{K_2^2-1})K_1 S}) \quad (6)$$

The maximum peak of the curve in Figure 2a can be calculated and is proposed by the authors in equation (7).

$$S_m = \frac{\ln(-K_2 - \sqrt{(K_2^2-1)}) - \ln(-K_2 + \sqrt{(K_2^2-1)})}{2K_1\sqrt{(K_2^2-1)}}$$

$$(7)$$

The shear strength of soil (τ) can be defined as the maximum, or limiting, value of shear stress that may be induced within its mass before the soil yields[3]. A Mohr diagram plotting ground pressure vs. shear stress, figure 3, shows the state of the stress for any orientation of a reference axis. The Mohr circle can only expand to a critical point before failure occurs. The line tangential to where failure occurs is the Mohr-Coulomb failure line.

145

The equation of this line is y = mx + b where b is the coefficient of cohesion, m is tan (ϕ), ϕ is the frictional angle, and x is normal stress or ground pressure. This line is the fundamental approximation to the maximum shearing strength, τ_m, of a particular type of soil and has been adopted as the definition of strength in land locomotion.
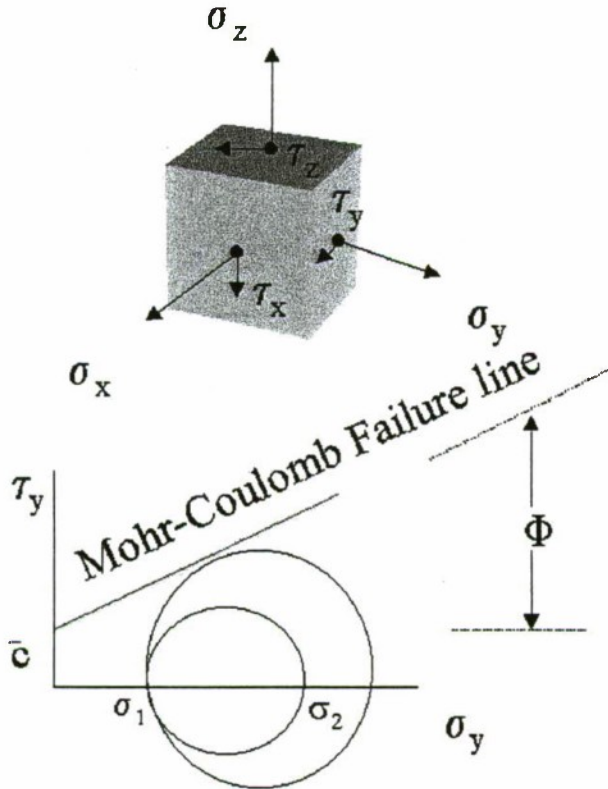
$$\tau_m = c + p\tan(\phi) \qquad (8)$$

Figure 3   Mohr-Coulomb Failure Line

Since the portion contained in brackets (Eq. 6) is dimensionless, the value of $K_3 / 2K_1\sqrt{(K_2^2-1)}$ must have the units of lb/in^2 and the value of K_3 may be expressed in the following manner.

$$K_3 = \frac{2K_1\sqrt{(K_2^2-1)}(c+p\tan\phi)}{[e^{(-K_2+\sqrt{(K_2^2-1)})K_1S} - e^{(-K_2-\sqrt{(K_2^2-1)})K_1S}]_{max}} \qquad (9)$$

Now equation (6) can be simplified:

$$\tau = \frac{(c+p\tan\phi)}{y_{max}}[e^{(-K_2+\sqrt{K_2^2-1})K_1S} - e^{(-K_2-\sqrt{K_2^2-1})K_1S}] \qquad (10)$$

where y_{max} is the largest value within the brackets. The slip distortion and the amount of slip are related. The distance of shear (S_m) is equal to the speed of the slip times the time in which it occurs.

$$S_m = v_S t \qquad (11)$$

However, the speed of slip is equal to the speed of the tire or track minus the actual speed:

$$v_S = v_t - v_a \qquad (12)$$

$$S_m = (v_t - v_a)t \qquad (13)$$

and $t = d/v_t$, where d is the distance where S_m has occurred.

$$S_m = d(1 - \frac{v_a}{v_t}) = i_o d \qquad (14)$$

The amount of soil distortion that takes place at any point at a distance x from the front of the ground contact area is equal to

$$S = S_m(x/d) \qquad (15)$$

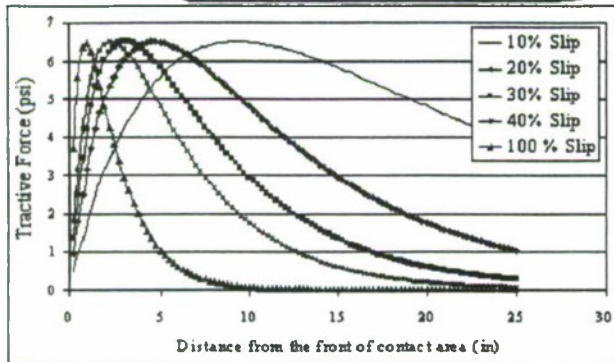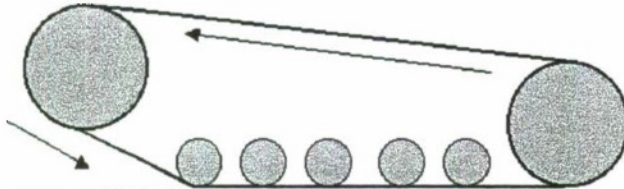So, $\qquad\qquad S = i_o x \qquad (16)$

Equation (16) then allows for a relationship between tractive force and slip. Figure 4 shows the shear force of a tracked vehicle in two types of soil. The top graph is of highly frictional undisturbed firm silt. At ten-percent slip, shear is produced along the entire track; but it is clear that the front half of the track is producing the most of the force[2]. As the vehicle begins to experience more slip, most all of the shearing force is produced at the front of the tracked vehicle. In fact, the back half of the track begins to produce no shear and actually increases the resistance by creating drag.

The second type of soil has a high cohesive property such as wet clay. At all values of slip, the entire length of the track is producing shear in relatively equal amounts along the length of the track.
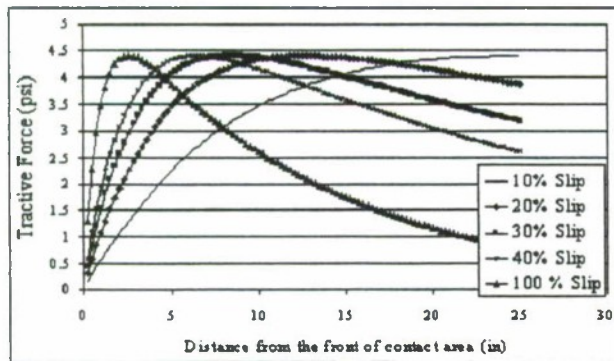
While in motion, a track or wheel develops a force produced by the shearing strength of soil. This force H is called the gross tractive effort or soil thrust. The tractive effort is the integral of the shear produced by a tire or track. By substitution of equation (10),

146

$$H = \int_0^d \tau_x \, dx \qquad (17)$$

$$H = \int_0^d \left\{ \frac{(c + p \tan \phi)}{y_{max}} \left(e^{(-K_2 + \sqrt{K_2^2 - 1})K_1 ix} \right. \right. \qquad (18)$$
$$\left. \left. - e^{(-K_2 - \sqrt{K_2^2 - 1})K_1 ix} \right\} dx \right.$$

Undisturbed Firm Silt

Undisturbed Settled Sandy Loam

Figure 4 Tractive force in different soil types

Figure 4 shows the soil distortion at any distance x from the front of the ground contact area. The top graph shows tractive effort produced in undisturbed firm silt. The maximum tractive effort is quickly produced a short distance from the front of the vehicle and the rest of the track produces very little even at a very low percentage of slippage. The bottom graph shows the same track moving in an undisturbed settled sandy loam.

It is often thought of the heavier a vehicle is the greater its tractive effort. Much experience gives credibility to this statement but is it valid for all soil types? In order to answer this question, consider equation (19). Soil thrust is defined as the addition of two different soil strengths. One is from frictional properties and the second is from its cohesive properties.

$$H = A \cdot c + W \cdot \tan \phi \qquad (19)$$

If a soil type such as dry sand is chosen, a homogenous sample would contain no cohesive properties, Therefore c=0, and equation (19) is reduced to W·tanφ. There is no question as the weight is increased the amount of soil thrust increases proportionally.

If the same vehicle is operated in a plastic soil such as saturated wet clay, the frictional component of the soil is equal to zero (φ=0). Equation (19) is reduced to $A \cdot c$ where A represents the contact surface area of the vehicle. A higher value of thrust is only obtained by an increase in contact surface area.

To answer the question in a more direct approach, vehicles that traverse in highly frictional soils benefit from an increase in payload. However, in soil types with high moisture contents or very cohesive, vehicles benefit by an increase in contact surface area. An increase in weight in this type of soil would be a liability[5].

BDTM

BDTM was established to give a first pass general evaluation of robotic vehicle mobility performance. It is a simple, linear one-degree of freedom (1-DOF) model that has been created in a spreadsheet format. The model assumes that the soil is homogenous and the loading effects on the soil are linear. A tracked vehicle and a wheeled vehicle can be simulated at one time. These vehicles are evaluated on their tractive force, tractive effort, soil sinkage, drawbar pull, and tractive coefficients (DP/W).

Inputs

The inputs into the program are divided into three categories. The first set of inputs are general vehicle information. These include the width and length of one track or wheel in contact with the ground. A corresponding code number relates to the actual shape of the print that the vehicle leaves on the ground. Other items include the number of tracks or wheels, contact area, and vehicular weight.

The second set of inputs describes the vehicle trafficability, or conversely, the vehicle performance in a given terrain. These define the strength, sinkage, and

147

slippage that a vehicle would experience in a specific homogenous soil type. Most of these parameters are obtained from the Bevameter, which is a device created by Bekker for this purpose[1]. These include the depth of the plate sinkage, the modulus of soil deformation in cohesional and frictional soil, the exponent of soil deformation, and the coefficients of slippage. A separate section in the program provide these for different types of soil. Other parameters such as the coefficient of cohesion and the angle of friction are calculated off the Mohr-Coulomb failure line.

The third set of inputs is used for the calculations of WES mobility indexes. The purpose of such inputs is to relate the Bevameter values to the WES cone index. This then allows for the comparison of results obtained through the NRMM mobility model. WES mobility indexes are defined by equation 20 and 21 [4]. The mobility index for a tracked vehicle is calculated by:

$$
MI = \left[\frac{\text{contact pressure x weight factor}}{\text{track x grouser factor factor}} + \text{bogie} - \text{clearance factor factor} \right]
$$
$$
\text{x engine x transmission} \quad (20)
$$
$$
\text{factor factor}
$$

and the mobility index for a wheeled vehicle:

$$
MI = \left[\frac{\text{contact pressure x weight factor} \quad \text{wheel}}{\text{Tire x grouser factor factor}} + \text{load} - \text{clearance factor factor} \right]
$$
$$
\text{x engine x transmission} \quad (21)
$$
$$
\text{factor factor}
$$

Outputs

The outputs are arranged into seven different categories. The first set is the theoretical soil thrust that the soil should support. This comes from the Mohr-Coulomb failure equation multiplied by contact area. It is expressed in equation (19) where $W \cdot \tan\phi$ is for the frictional composition of the soil and $A \cdot c$ is due from cohesion. Since most soil is a mixture of these two compositions, soil thrust in average soil is from the addition of these two terms.

The next output set is for strengths and pressures. The normal force exerted on the soil is due to loading from the vehicle and is referred to as the ground pressure. The maximum soil strength is Mohr-Coulomb failure equation calculated at the corresponding ground pressure. From the

stress-strain curves of soil, Bekker noticed that they are identical to the displacement (x) and natural time frequency (ωt) of an aperiodic vibration[2]. The equation for tractive force was derived from this remark and is shown in equation (10). Its soil properties and the amount of slip distortion evaluate the tractive force. This is a product of the distance from the front of the track multiplied by the percentage of slippage the vehicle is experiencing. Equation (18) expresses the tractive effort in terms of soil properties, contact area, load, and slip for a given type of soil defined by its K_1 and K_2 constants.

To evaluate sinkage in frictional and cohesive soil, Bekker derived a formula from his Bevameter

$$
z = \left[\frac{p}{k_c / b + k_\phi} \right]^{1/n} \quad (22)
$$

where p is the ground pressure, b is the width of the track or tire, k_c and k_ϕ are frictional and cohesive modulus of soil deformation, and n is the exponent of soil deformation. This equation answers why wider tracks or tires on vehicles with the same ground pressure sink deeper.

Not all soil thrust can be accounted for the production of useful work. Some of the soil thrust is lost in the form of energy. The energy loss that compose the external resistances are caused by compaction of soil, bulldozing, and dragging. It has been shown that the portion wasted for overcoming compaction resistance may be expressed by

$$
R_c = \left(\frac{1}{(n+1)(k_c + bk_\phi)} \right) \left[\frac{W}{l} \right]^{\frac{n+1}{n}} \quad (23)
$$

where W is weight in pounds and l is the length of the tire or track in contact with the ground. It can be noted that from equation (23), the longer the contact area the smaller the compaction resistance. Bulldozing is the visible pushing of soil mass in front of a vehicle. For this model the resistances that are due from bulldozing are neglected. Also the resistances that occur from trapping the soil and dragging it are neglected.

The drawbar pull (DP) is the total thrust minus the total resistances. It is customary to view the difference as the vehicle's ability to move. If the total is zero or negative, then the locomotion of the vehicle will stop. In BDTM, there are three different values of DP. The first is considering soil thrust developed purely off of soil

parameters. The second DP value is including the additional thrust that is created by the action of grousers or treads. The Mohr-Coulomb line equation is then modified for this result:

$$H = blc(1 + 2h/b) + W \tan\phi\left(1 + 0.64\left[(h/b)\cot^{-1}(h/b)\right]\right)$$

(24)

where b is the width, l is the length, h is the height of the grouser or tire tread, c is the coefficient of cohesion, and ϕ is the angle of friction. The last value of DP that is in BDTM is the value of the total tractive force evaluated at a certain slippage at a specific distance from the front of the contact area. A common comparison used to evaluate vehicles is to normalize these DP by there weight. This is often called the traction coefficient and should not be used as a stand-alone measure in evaluating vehicles.

The final set of outputs is devoted to cone index (CI) and mobility index (MI) conversions. Waterways Experiment Station (WES) came up with a way to measure soil parameters. The cone index is the parameter that is obtained by using their cone pentrometer device. The CI values are obtained by converting Bevameter values into CI values from equation (25)[6]. The conversion was proposed by Janosi and tested by WES in 1964. It was shown to be consistent within the limits of accuracy.

$$CI = 1.625\left(\frac{\frac{k_c}{(n+1)}\left((z+1.5)^{n+1} - z^{n+1}\right)}{+0.517k_\phi\left(\frac{(z+1.5)^{n+2}}{(n+1)(n+2)} + \frac{z^{n+2}}{n+2} - \frac{(z+1.5)z^{n+1}}{n+1}\right)}\right)$$

(25)

Charts

In BDTM, there are four charts that provide useful information. The first chart is traction coefficients versus k values. This provides curves for both the tracked and wheeled vehicle for the traction coefficient in different strengths of soil. Figure 5 shows a tracked vehicle and the same vehicle with tires in a mostly frictional soil type. It can be seen that the tracked vehicle can easily traverse soil with less consistency than the same vehicle with tires on.
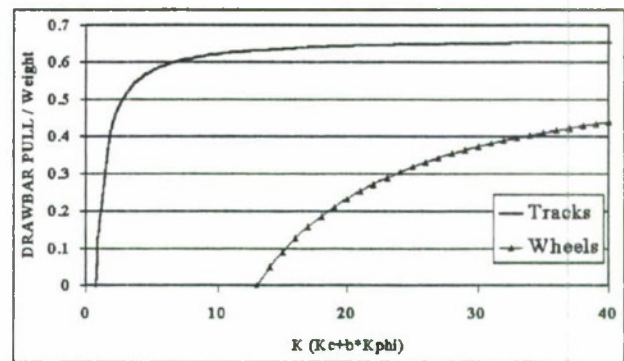

Figure 5 DP/W vs. K values

The next two charts show the tractive force produced under the contact area of the track or tire.

These curves are made at 10, 20, 30, 40, and 100 percent slip. Figure 6 shows a tracked vehicle in an undisturbed settled sandy loam. It is shown that the track produces force constantly down the contact area of the track. Even at various levels of slip.
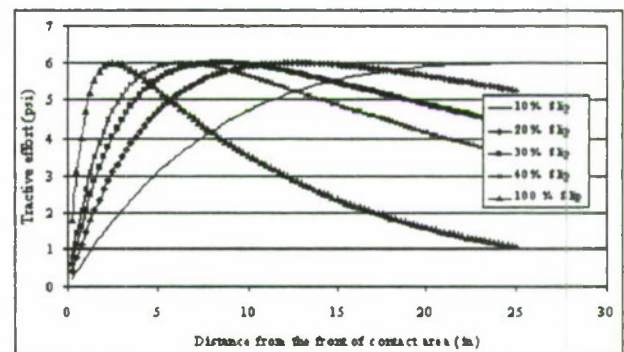

Figure 6 Tractive Force vs. Distance Under Track

The last chart displays tractive effort per unit area with soil distortion. The amount of work that is accomplished as the amount of soil distortion occurs. This is evaluated as the slippage increases. Figure 7 shows tractive effort versus slip at three different soil types. Soil type A is a highly frictional soil type and can be seen that almost all the tractive effort is produced when the vehicle experiences less than ten percent slippage. On the other-hand, soil type C is a plastic or cohesive soil type. Tractive effort is produced relatively uniform regardless of the soil distortion or slippage experienced.

149
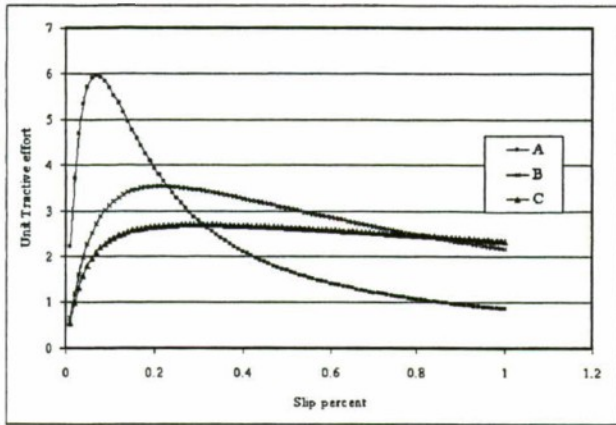
A. Cohesionless or loose frictional soil
B. Mixture of A and C
C. Cohesive or plastic soil
Figure 7 Tractive Effort vs. Slip

EXAMPLES

Two examples have been provided to demonstrate the ability of the model.

Example 1

The track vs. tire case has been argued quite extensively. It has been slated by some that a low-pressure pneumatic tire can perform as well as a track. What does the BDTM predict for small robotic platforms?

To address this question, lets look at an example of a robotic vehicle traversing in a highly cohesive soil type such as wet clay. A small four-wheeled robotic platform with 12" diameter tires as shown in figure 8a leaves a rectangular print 3"x4". The weight of the platform is 1000 lb. and the tires are located a distance of 36" apart. Each tire has a total contact area of 12" and an overall surface contact area of 48". The ground pressure of the vehicle is 21 psi.

The model shows the vehicle sinks to a level of 2.3". At this depth, the resistance to motion created by compacting and bulldozing is greater than the maximum soil thrust generated. The Drawbar-Pull is a negative value indicating that the vehicle is incapable of moving.

It is often though that an increase in payload could help in this situation. When a 200-lb payload is added to the robotic vehicle, the vehicle begins to sink deeper. The resistance to motion increases. Drawbar-Pull remains a negative number and the vehicle still is incapable of moving.

When the diameter of the tire is increased to allow for a 3"x6" print as shown in figure 8b, the total surface area is increased to 72" squared and the vehicle only sinks to a level of 1.5". The amount of resistance to motion has decreased to a level that the vehicle is capable of moving. This is indicated by a positive Drawbar-Pull; however, the amount of DP that is produced is minimal.

It is seen that an increase in tire diameter, which is an increase in surface contact area, leads to an increase in DP. The next logical step would then be to continually increase the diameter of the tire until the desired amount of DP is obtained. This approach leads to other problems such as turning radius and for our purpose is not practical.

A possible solution to this is to add another set of wheels. Figure 8d displays the robotic vehicle with six wheels. When the 12" diameter wheels are used leaving a 3"x4" print, the amount of surface contact area is equivalent to the four-wheeled vehicle with enlarged tires. Therefore, the same results may be obtained by using six smaller wheels than with four enlarged. By increasing each of the six tires to allow for a 3"x6" print for each, the total surface area is increased to 108" squared. The vehicle's ground pressure has decreased to 9.1 psi and sinks 1" in the ground. The amount of Drawbar-Pull doubled.
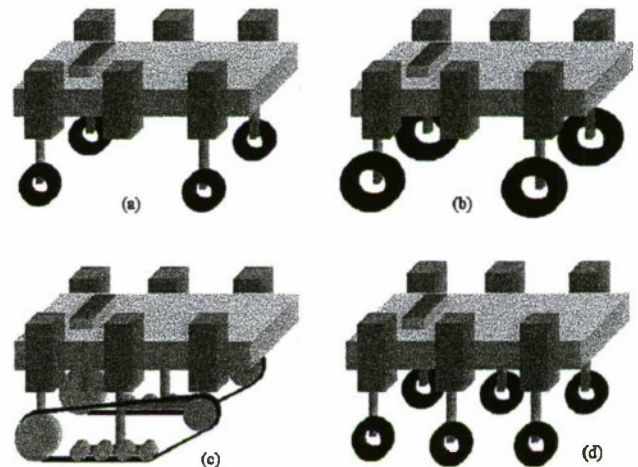

Figure 8 Robotic Vehicles

When the vehicle is outfitted with a track that is 25" long and 3" wide as shown in Figure 8c. The ground pressure has decreased to a level of 6 psi and sinks .6" in the soil. The drawbar pull is 14 times greater than the 4-wheeled vehicle, 4 times greater than the vehicle with enlarged tires, 4 times greater than the six-wheeled vehicle with 12" diameter wheels, and 2 times greater than the six-wheeled vehicle with enlarged wheels. When an additional 200-lb payload is placed on the tracked vehicle,

the tractive effort remains unchanged, and the DP actually decreases due to sinkage.

This example shows that in a plastic soil type such as wet clay. The vehicles ability to traverse is dependent on the amount of contact surface area.

Example 2

For this example, we will look at the question in example 1 about track vs. tire but in a highly frictional type of soil such as dry sand. The same robotic vehicle platforms as shown in Figure 8 has been selected.

When the vehicle operates with tires that leave a 3"x4" rectangular print, the tractive force and soil thrust produced are very comparable to the vehicle outfitted with a 3"x25" track. The track outperforms the tire only 1.5 to 1. If the 4-wheeled vehicle is outfitted with the oversized tires leaving a 3"x6" surface contact print. The ratio is decreased to 1.2 to 1. The six-wheeled vehicle with the 3"6" print tires are almost 1 to 1.

It is quite interesting to note that when the vehicles are experiencing more slip. The 4-wheeled vehicles actually start to outperform the tracked vehicle. This begins to occur at around 33% slip for the 3"x4" print and 24% for the oversized tire.

Another thing that is fascinating is when the payload is increased for the tracked vehicle; the tractive force, soils thrust and drawbar increased respectively. When the 4-wheeled vehicle payload increased, the tractive force and soil thrust increased; but the drawbar pull decreased. By decreasing the weight of the 4-wheeled vehicles by 200-lb, the track only outperformed by 1.3 to 1 for the 3"x4" print and 1.1 to 1 for the oversized tire.

It can clearly be seen by this example that a lower weighted-wheeled vehicle can perform as well if not better than a tracked vehicle in highly frictional soil types such as dry sand.

SUMMARY

The BDTM model is our first attempt at developing a spread sheet formulation of Bekker's early Terramechanics model for off-road vehicles. The fundamental assumptions in Bekker's work are that shear/slip is the primary mechanism for generating tractive vehicle forces, a 1-DOF over-damped, linear model describes vehicle motion, and that soils plastically deform under ground pressure forces. A total of seven parameters constrain the model to a unique set of vehicle/terrain characteristics. The BDTM assumes a generic soil condition, which is a combination of frictional (sandy) and cohesive (clay) type soils.

A primary objective of this paper is to evaluate Bekker's formalism as a tool to perform mobility tradeoffs between wheeled and track vehicles. BDTM is able to accurately predict the percent slip and drawbar-pull characteristics of track vehicles for several different soil conditions. In addition the model predicts soil distortion at any x from the front of the ground contact area, maximum tractive effort, and soil sinkage for various soil conditions. Several examples compared the relative mobility advantages of wheel and track vehicles for different types of soil conditions. In some cases 4-wheeled vehicles begin to outperform track vehicles for high slip conditions.

Future analysis will compare six and eight wheel vehicles with track vehicles. The former should have considerably better mobility performance than their 4-wheel counterparts. The wheel vs. track comparisons will be made for different weight classes as well ranging from less than 100 pounds to the M1 tank. In addition the link between BDTM and NRMM will be further developed by relating the mobility and vehicle cone indicies to Bekker's model.

REFERENCES

[1] Bekker, M.G. 1960 *Off-the-Road Locomotion* University of Michigan Press, USA

[2] Bekker, M.G. June 1955 *"Wheels or Tracks"* Automobile Engineer

[3] Whitlow R 1990 *Basic Soil Mechanics* (2nd Ed) Longman Group, UK p207

[4] Wong, J.Y. 1989 Elsevier Science Publisher B.V.,Netherlands

[5] Bekker, M.G. Land Locomotion Research Report 13, 1957 *"Terrain Evaluation in Automotive Off-the-Road Operations"* OTAC, Detroit

[6] Bekker, M.G. 1969 *Introduction to Terrain-Vehicle Systems* University of Michigan Press, USA

SAM Signatures from the Chicken Little Vehicular Fleet for Use in Modeling Techniques

John Sledge
46OG/OGML
CHICKEN LITTLE Program Office
Eglin AFB, Florida 32542
August 1999

ABSTRACT

Ground vehicles are emitters of seismic, acoustic, and magnetic (SAM) signatures of military interest. These sensors are low cost and passive devices and are used for the gathering of signatures from ground vehicles. Vehicular scenarios can include from the benign to warhead detonations in the vicinity of the target. The signatures are of a robust nature and lend themselves to be exploited for seismic, acoustic, and magnetic models. The paper begins by defining the procedures and standards for the collection of quality signatures. Identification and discussion of SAM vehicular sources and battle damage assessment (BDA) technique follows this. Acoustic propagation models and signature data basing discussions related to the SAM data are essential components. Finally a composite of seismic, acoustic, and magnetic signatures provide a realistic picture of what can be accomplished.

INTRODUCTION

The use of laboratory grade seismic, acoustic, and magnetic devices for the gathering of quality signatures is essential. These sensors which operate in the passive mode, have been used for military purposes since the Vietnam War. Such sensors are small, low priced, and can consume low power levels during operation. Additionally, these sensors are very adaptable to harsh environments.

Vehicles of military interest are emitters of seismic, acoustic and magnetic (SAM) signatures. The signatures when properly collected and correlated can provide a fairly accurate description of the source. Things such as types of engines, number of engine cylinders, type of exhaust, track or wheel noise, sound intensity, track frequency, turbine engines, turbo charged engines, on-board and induced magnetic fields, are some of the vehicular sources that are sensed by SAM devices. These signatures are frequently used for the exploitation, detection, and classification of military vehicle systems.

The immediate vehicular background environment activity also provides a wealth of information that contributes to the overall signature. These are the weather, battlefield activity, explosions, radio communications, radar emissions, and other man made and natural activities.

Data basing of the seismic, acoustic, and magnetic signatures plays an important role in the SAM model development. It is important to insure that the data collection cycle, calibration to Range Commanders Council (RCC) Standards, and documentation are part of a quality model.

STANDARDS, SENSOR DEPLOYMENT, and DATA COLLECTION

Data Collection Discipline

Seismic, acoustic, and magnetic data collection standards are the key to quality (SAM) signatures. Standards are necessary within the Department of Defense (DOD) to control the constraints of variable site conditions on SAM sensor emplacement, deployment, and performance. These standards are used to establish compatibility between the sites at DOD centers. Regardless of where the test is conducted, the same SAM standards are employed.

The RCC Standards provide the discipline for the collection of signatures. Included in the standard is the documentation of sensor parameters with definitions, instrumentation procedure, ground truth guidelines,

calibration technique, and the meteorological requirements, and other documentation required.

Acoustic Sensor Characteristics and Architecture

Acoustic sensors are used to convert pressure variations in the atmosphere to electrical signals. The simplest microphone deployment is a single unit. Other arrays include circles, triangles, linear designs, and underground. Figure 1. Shows a triangular array of ground placed microphones.

Figure 1. Shows a Ground Placed Triangular Acoustic Array

Acoustic arrays typically used to provide bearing estimates to a sound source and for removing the effect of a nuisance sound source. Placement of acoustic sensors is very important. In some instances, the sensor must be isolated from the ground surface to avoid the introduction of seismic vibrations. One technique for isolating the microphone from the ground is to employ a small-inflated inner tube between them. Figure 2. Shows this type of placement.

Figure 2. Single Microphone with Isolating Inner Tube

Another form of acoustic linear array is the suspended acoustic array. The single-sided array consists of nine microphones spaced at 10-degree intervals. This type of array provides a top-down perspective to an acoustic source. Another type of suspended acoustic array is the double-sided array. This device consists of eighteen microphones spaced at 10-degree intervals. Both arrays are used to gather acoustic data samples to simulate air to ground type munitions that have acoustic sensors. The suspended acoustic arrays are also used to collect data samples to simulate an AUV acoustic sensing gathering activity. Figure 3. Shows the single sided acoustic array.

Acoustic data samples from the suspended acoustic array will be discussed later in this paper.
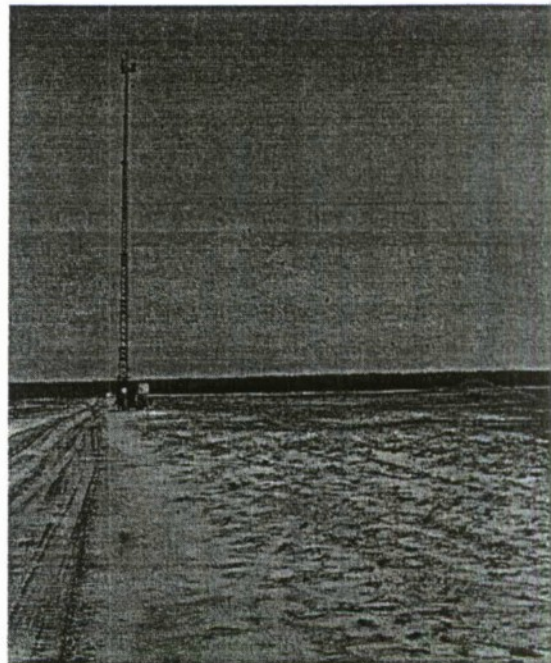

Figure 3. Single Sided Suspended Acoustic Array

Seismic Sensor Characteristics and Deployment

The geophone is a relatively simple device that employs a magnetic mass in a coil to produce a signal. A geophone that is implanted into the ground vibrates with the ground, and the coil moves with the supporting case while the magnetic mass remains stationary. The motion between the mass and the coil generates an electrical signal, making the geophone a self-generating device. There are two basic types of geophones, those that measure motion on a single axis or on three axes.

The simplest form of a seismic sensor is the single axis geophone. This device measures vertical motion and is

153

most frequently used for collecting seismic signatures. The 3-axis geophone provides a broader set of seismic signatures. Seismic data from the 3-axis geophone provides information that enables the user to analyze the propagation of seismic waves (how the different types of seismic waves compose the measured signal). A 3-axis seismic sensor is recommended when collecting seismic signatures.

Geophones may be planted on the surface of the ground or buried at some depth, depending on the application. Regardless of the type of geophone or the emplacement, there is a need to insure that the geophone is in good contact with the soil. It is recommended that the RCC Seismic Standards be employed when planting geophones. Figure 4. Shows a 3-axis geophone collocated with an acoustic and magnetic sensor.
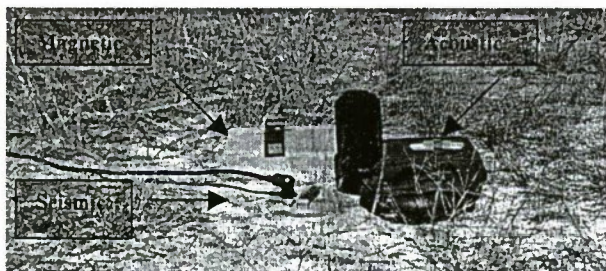

Figure 4. Seismic Sensor Collocated with Acoustic and Magnetic Devices

Magnetic Sensor Characteristics and Placement

The magnetic sensor consists of a 3-axis flux-gate magnetometer. This device is used to make precision measurements of magnetic fields with a sensitivity level of plus or minus one (1.0) Gauss. The magnetometer measures the magnetic field sensed by each of the orthogonal axes. Its signal output is linear with applied field over the range of plus or minus1000 mGauss.

Ground vehicles are composed of ferromagnetic materials and when moving in a uniform magnetic field, they produce magnetic field variations, which can be detected by a magnetometer. The magnetic field variations are produced by a magnetic dipole moment and are functions of the distance between the magnetic sensor and the target moment.

The calibration of the magnetometer is conducted at a facility where documentation is traceable to National Institute of Standards and technology, NIST.

Placement of the magnetometer sensor in the field is accomplished by isolating the sensor from the ground surface. This is accomplished by using a piece of concrete, like a stepping stone, for isolation purposes. Figure 4. Shows a magnetometer sensor in the field.

Other applications require that the magnetometer be mounted in an elevated position. This type of setting is used for those tests that require overhead vehicular measurements. Another area of magnetic interest, is the measurement of the on-board moment. These are the signatures that emanate from vehicles that have for their source active electrical systems.

VEHICULAR SAM SIGNATURE SOURCES

Sound sources including all types of machinery and human activity has its own peculiar characteristics. Quite enough the sound is distinctive to allow identification and tracking of the source. Away from the source, the sound generated is modified by atmospheric propagation which depends on a variety of meteorological and surface terrain conditions. A careful characterization of a source or a receiving system must account for the propagation conditions present as well as characteristics of the measurement and recording system.

When planning acoustic measurements in the atmosphere, many variables must be controlled or documented. At a minimum these include the characteristics of the sound source, a description of the receivers and recording system, characterization of meteorological conditions, and the ground along the propagation path, a description of the terrain, and notes on background or nuisance sound sources.

Despite modern engineering designs to limit seismic, acoustic, and magnetic emanations produced by military vehicles, SAM sensors and new technology can still detect and classify these vehicles.

Atmospheric Acoustic Signatures

The primary source of acoustic noise in a military vehicle is that produced by the engine and it is delivered through the exhaust. Other sources include, auxiliary power units, fans, transmissions, differentials, pony engines, tires, and tracks to mention a few.

Earlier it was mentioned that the vehicle engine exhaust system was the primary noise source. Figure 5. Shows spectrogram from a 12-cylinder diesel engine. Closer

inspection of the spectrogram indicates that numerous engine harmonics populate a certain spectral region. Two harmonics dominate the spectral region of interest. These are the 6th and the 12th harmonics. This is a typical feature exhibited by V-12 engines that have a dual exhaust system. Additionally, the presence of tire noise is evident as the vehicle moves. Gear changes that took place before the vehicle attained a specific speed are shown.
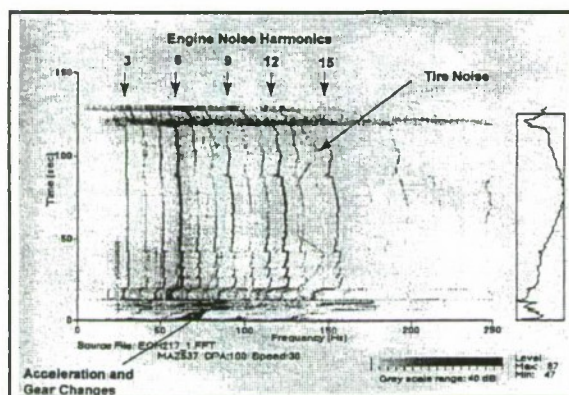

Figure 5. Acoustic Spectrogram of Military Wheeled Vehicle in Motion

Seismic Signatures

Seismic signatures are those generated by means of direct ground contact by an area of the vehicle. A good source of seismic energy is the track on armored vehicles. This seismic source produces a track frequency as a result of vehicle travel. Wheels from military vehicles are another source of seismic energy.

The seismic spectrogram shown in Figure 6 shows a wealth of information about the military wheeled vehicle as it undergoes a dynamic routine. Suspension vibrations, tire noise, and an indication of acoustic coupling make up the seismic spectrogram. The acoustic coupling can be traced back to the engine acoustic exhaust harmonics. Yes, geophones can pick-up acoustic engine harmonics. When collecting seismic signatures it is necessary to collocate a microphone with the geophone. This is necessary to compare and determine the origin of the signatures. The Seismic RCC Standard describes this technique in detail.
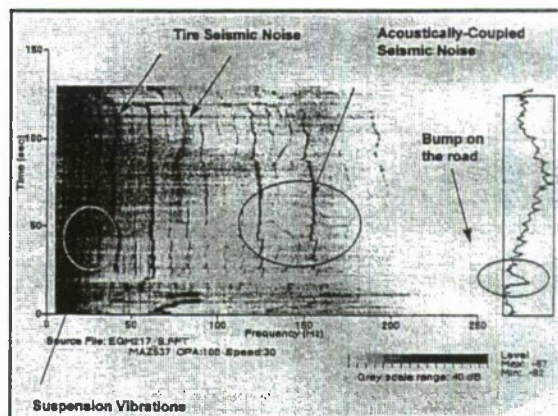

Figure 6. Seismic Spectrogram of the Military Wheeled Vehicle in Motion Depicted in Figure 5

Magnetic Signatures

Magnetic sources are divided into the induced moment and the on-board moment. The induced moment source is that produced by the mass of the vehicle as it moves in the vicinity of the magnetometer. The on-board moment source is produced by all the other electrical systems that are part of the vehicle. These include electric motors, centrifugal oil filters, radar emissions, tachometers, radios, etc. Figure 7 is a good example of a magnetic signature that depicts both magnetic moments.
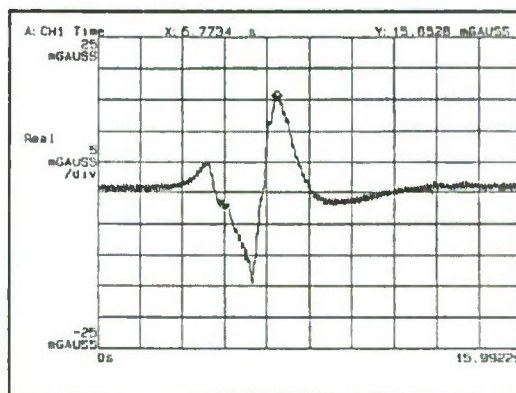

Figure 7. Magnetic Signature depicting the Induced and On-Board Moments

The magnetic signature depicted in Figure 7 is a good representation that contains the induced and on-board moments. The sine wave is indicative of the induced moment and the modulation is contributed by the on-board moment.

SAM SIGNATURES FOR BDA

The Chicken Little Joint Program Office has pioneered the use of SAM sensors for battle damage assessment purposes. A recent test was conducted in cooperation with the German Government at the Meppen Test Facility, Meppen, Germany. The test involved the use of an operating SA-6 Radar vehicle and the complementary units typically found in this type of setting, and the detonation of an Improved HARM warhead.

BDA Acoustic Findings

An Improved HARM Warhead was detonated in the immediate area to determine the effects produced by the warhead. SAM sensors were deployed in the field to collect signatures from the explosion. Some of the SAM sensors were deployed one meter below the ground surface, and others remained on the surface. The sensors were placed 150 meters from the operating SA-6 Radar vehicle. In a situation like this, prior to any data analysis, it is advisable to look at the analog data stream in the time domain. A time history places the events in perspective. Figure 8. Shows acoustic data from a microphone buried one meter below the ground surface.
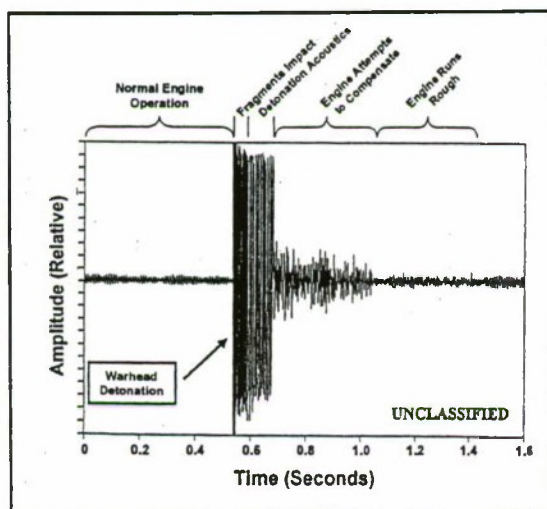

Figure 8. Acoustic Signal Time History from SA-6 Radar Vehicle Engine and Improved HARM Warhead Detonation

The dominant acoustic harmonics of an operating engine emanates from its exhaust port. This acoustic product, composed of low frequencies, propagates over long distances. The SA-6 Radar vehicle followed this trend. Although the APU was on at the same time, its acoustic product did not propagate over a long distance. The environment readily attenuates acoustic high frequencies from the APU.

The engine exhaust acoustic product is a combination of harmonics. These harmonics can be found by doing a spectrogram of the analog signal. The dominant harmonic, or engine firing frequency, results from the engine operating at a specific rpm value.

In Figure 8, the SA-6 Radar vehicle is stationary and the engine is operating at a specific rpm. This is defined as normal engine operation. The warhead detonates and specific areas of the explosive signature become evident. The engine trying to compensate immediately follows the explosion event. It is evident that loads are being placed on the engine. Engine loading is the result of breakers and fuses from the electrical system blowing and popping. Warhead fragments have hit the active SA-6 Radar dish causing electrical shorts. Afterwards, the engine stops compensating because all electrical malfunctions are no longer taking place. The radar is now off line. Now the acoustic time line history plot shows that the engine runs rough. Actually, the rpm value changed as result of the detonation.

BDA Magnetic Findings

A military vehicle such as the SA-6 Radar is a self-propelled ferromagnetic mass. It is therefore reasonable to expect that the detection and classification scheme based on a stationary magnetic sensor would prove useful in situations involving vehicles of military interest. The physical principal involved is that a ferromagnetic mass exhibits an induced magnetic moment when placed in the earth's magnetic field. The magnetic field of this induced dipole distorts the uniform geomagnetic field that can be detected by a magnetometer that produces a signature depicting the disturbance.

In addition to the induced moment, there exists a second magnetic moment, which contributes to a vehicle's magnetic signature. This moment, called the "on-board" moment, has for its origin electrical phenomena and mechanical motion created by the vehicle. A portion of the on-board moment is also due to the residual magnetic effects in the vehicle material due to long-term positioning in the geomagnetic field.

BDA Magnetic Findings

The stationary SA-6 Radar vehicle was located 150 meters away from the magnetometer. In this configuration, the magnetometer has to rely on sampling

156

"on-board" moment activity as it takes place aboard the SA-6 Radar. The primary source is the on-board magnetic activity that results from the one (1) Watt radar magnetic emanations. Sampling of the induced moment at this distance is a problem because the data would be lost in the background noise. Thus, all efforts must be focused on the on-board moment activity.

The magnetometer was set up to sample emanations in a spectral region of less than (forty) 40 Hz. Magnetic data will now be presented showing two areas of interest. The magnetic samples before and after the warhead detonation. All magnetic data samples are presented in the time domain.

Magnetic Background Noise at the Meppen Site

A sampling of magnetic background noise was acquired to establish a baseline. This sample was collected with the SA-6 Radar vehicle stationary and all systems off. Figure 9 shows that the average magnetic background was 15.0 mGauss.
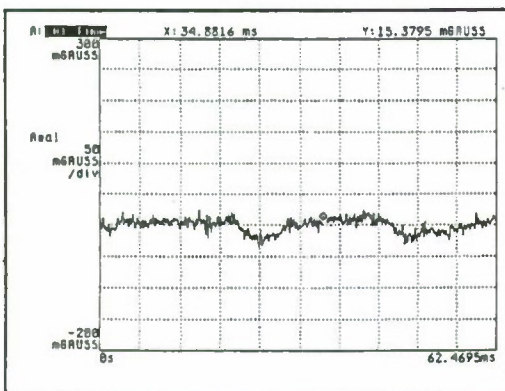

Figure 9. Magnetic Background Noise at the Meppen Test Site, Germany

The SA-6 Radar vehicle engine is now in operation and the radar is emanating 1 Watt of power. It must be noted that the magnetometer is sampling the backside of the SA-6 Radar dish. Figure 10 shows that the magnetic levels now average 140.0 mGauss.
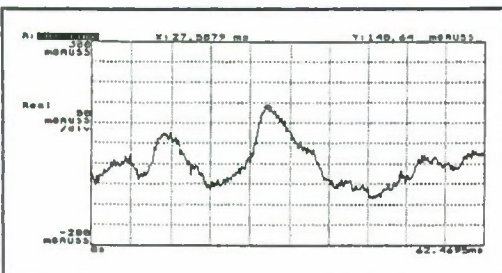

Figure 10. Magnetic Signature from the Backside of the SA-6 Radar Vehicle, Before Warhead Detonation

The HARM warhead detonation has taken place. Magnetic effects from the warhead detonation are added to the radar's magnetic field. Figure 10 shows the magnetic data sample immediately after the detonation.
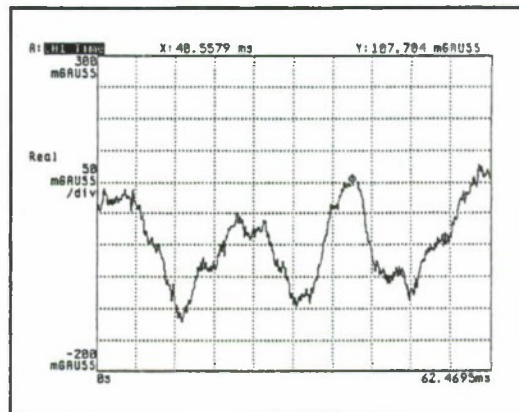

Figure 10. Magnetic Data Sample from the Backside of the SA-6 Radar Vehicle, Immediately after Warhead Detonation

The radar begins to show additional signs of magnetic decay. Emanations begin to decrease rapidly. The signature begins to show random attempts by the vehicle electrical systems to compensate. Metal debris from the HARM warhead has penetrated areas of the radar dish. Other fragments found their way to critical sections of waveguide that connects the transmitter to the dish. See Figure 11. This data sample was collected 11 seconds following the detonation.
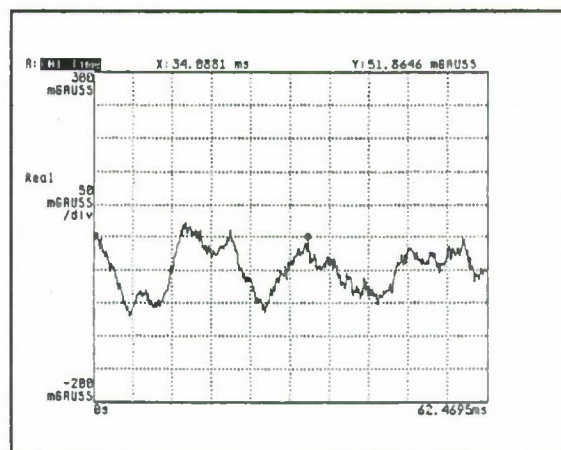

Figure 11. Magnetic Signature from the Backside of the SA-6 Radar Vehicle, Eleven Seconds After Warhead Detonation

The radar continues to show significant signs of magnetic decay. Emanations are now at a very low level. The

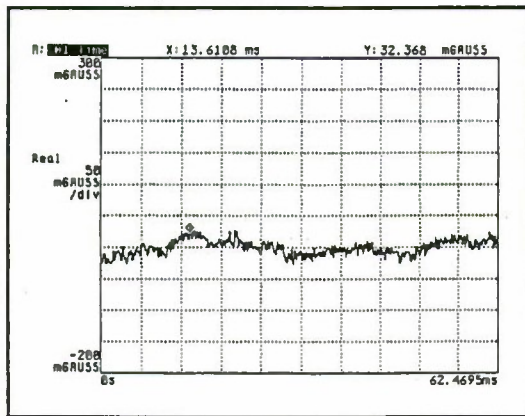signature continues begin to show a random effect. See Figure 12.

Figure 12. Magnetic Signature from the Backside of the SA-6 Radar Vehicle, Thirty Seven Seconds After Warhead Detonation

The magnetic signature from the SA-6 Radar on-board moment is now approaching the background noise baseline value of 15.0 mGauss.

Magnetic Data Comments

It can be concluded that the magnetic sensors could be placed at a farther distance, if the SA-6 Radar Vehicle transmitter were to operate without a dummy load. The low frequency magnetic sampling of the on-board moment data from active radio frequency transmitters opens many areas of interest.

METEOROLOGICAL AND ACOUSTIC PROPAGATION MODELS

The environment has a dramatic effect on the transmission of acoustic signals through the atmosphere. Transmission affects different frequencies differently so the received signal spectrum also depends upon atmospheric conditions along the propagation path. The phenomena that most seriously affect the received signal are refraction, which results from gradients in wind or temperature, and scattering by atmospheric turbulence. During the day, the ground heating can result in high temperatures near the ground decreasing with height. At night, this situation often reverses. The speed of sound is proportional to the square root of temperature so temperature gradients result in speed of sound gradients. Since refraction depends upon a gradient in the speed of sound, the temperature and wind speed and direction should be measured as a function of height. As a rule of thumb, the measurement of temperature and wind should extend to a height approximately equal to one tenth of the longest propagation distance.

A minimal measurement system places temperature and wind sensors at different elevations on a tower. These are at surface, 1 meter, 3 meters, 5 meters, and 10 meters. For acoustic collections involving propagation distances greater than 100 meters the use of radiosondes is recommended. This methodology complies with RCC Meteorological Standards.

Data from these meteorological conditions and devices are used to develop propagation models.

SAM DATA BASE

SAM signatures should be archived in a database that is user friendly and can be accessed by an ordinary computer. The SAM data must be accompanied by ground truth. Ground truth is additional technical information that adds substance to the interpretation of SAM data. Documentation is the key factor. The SAM RCC Standards document ground truth requirements.

SUMMARY

Seismic, acoustic, and magnetic sensors are passive, non-line of site devices that have a place in the detection and classification of military vehicles. The uses of new signal processing methods with the SAM data have proved that these low cost devices have a place in the battlefield. SAM sensors are frequently used to monitor front line activity and gather intelligence regarding a threat.

The use of the SAM sensors for BDA applications at the Meppen facility yielded a wealth of information. SAM sensors were placed in the near and far field for the purpose of gathering vehicle information. The vehicle of interest was the SA-6 Radar. Geographical placement of the SA-6 Radar and the SAM sensors, placed the far field sensors at a disadvantage. Despite the impediments, the SAM findings were very encouraging in determining the health of a damaged vehicle.

Discussion herein included descriptive information about vehicular sources that can be sampled by SAM sensors. Samples of vehicle signatures show areas that are typically used in the business for exploitation purposes. Procedures for collecting signatures, calibration techniques, and meteorological data are dictated by Range Commander Standards. Data basing of quality SAM

signatures and updating the system is necessary. These inputs play a key role in the development of realistic SAM models.

REFERENCES

Range Commanders Council, SAM Signatures Measuring Standards Group, April 1999. *Seismics Signature & Methodology Standard,* Revision NGIC-4-14-99.

Range Commanders Council, SAM Signatures Measuring Standards Group, April 1999. *Architecture, Calibration Techniques, & Ground Truth, for the Collection of Acoustic Data in the Atmosphere,* Revision NGIC-4-14-99.

John Sledge, April 1999. *Seismic, Acoustic, and Magnetic Results from a US – German Test*

NATO, ET-08 Conference Committee, June 1999. *Data Basing of SAM Signatures.* ARL Technical Presentation

Jerry Prado, Sentech, Inc., June 1998, *Personal Notes Concerning Seismic and Acoustic spectrogram Techniques*

Yie Ming Chen, Northrop-Grumman, December 1997. *Acoustic Propagation Comments*

Using the TARDEC Acoustic Ground Array to Determine the Characteristics of the Band Track

Edward Shalis and Douglas Freese
U.S. Army Tank-Automotive and Armaments Command
Warren, MI 48090 USA

ABSTRACT

An acoustic ground array was initiated by SenTech under a Phase 1 SBIR (Ref 1) to enable TARDEC to measure and locate acoustic sources. This program was enhanced by a joint TARDEC-SenTech effort that concluded with an operational system. During August of 1998, the ground array was used to determine the expected noise/vibration reduction of the flexible composite band track.

During detection, detections that were within the vehicle's path were added and a single percentage point of detection was formed for that speed and range. These summations were repeated for other ranges so that percentages of detection could be plotted as a function of range, for several vehicle speeds. This was done for the T-130 steel track and the band track. Finally, differences between the two tracks were plotted as a function of range for several speeds. Results show that a significant noise/vibration reduction is offered by the increased compliance and flexibility of the band track. However, this is offset by the inclusion of cleats, which are required to provide the necessary traction. Chordal action is not eliminated since the cleats themselves form a rigid

series of "track shoes". The existence of the band track cleats is a limit to the further reduction of noise/vibration offered by the present band track.

BACKGROUND

The Standard Steel Track

Large vibrations originate in regular steel tracks when track shoes interact with suspension wheels and produce reaction forces in the hull that are transferred throughout the entire vehicle by the vibration field. A regular track, because of the finite size of its rigid track shoes, moves over the idler wheel and interacts with the sprocket in a series of chord-like-segments, instead of a continuous elliptical loop. This chordal action produces track shoe velocity variations, which lead to fluctuations in the track tension. The result is that the track tension changes produce forces at the sprocket and the idler wheels and the reaction forces then vibrate the hull. The track shoe chordal action (Ref 2) occurs at the track laying frequency, so that major portions of the hull vibrate at a fundamental frequency. The track laying fundamental frequency is determined by the vehicle's speed and the length of the track shoes and can be expressed in the following units:

Frequency (Hz) = 17.6*Vehicle Speed V (mph) / Shoe Length (in)

Equation (1)

For the baseline M113 vehicle, the track shoe length is 6 inches and the fundamental track laying frequency can then be approximately expressed as:

Frequency (Hz) = 3 * V (mph)

Equation (2)

The track input forces also contain energy at many other frequencies which produce a broadband of vibrational inputs into the vehicle hull and a broadband of frequencies in the radiated acoustic spectrum (Ref 3). Thus machine detectors, such as acoustic ground arrays that we will use to measure the track signature, detect a fundamental frequency and its harmonic line set of track laying tonals superimposed on a broad band track signature.

Steel Track Acoustic Spectrograms
Figure 1 shows an acoustic spectrogram of the baseline vehicle moving at 10 mph while measured by a ground array located 900 feet from the vehicle. The fundamental track laying frequency of the baseline track occurs at 30 Hz, which is consistent with equation 2. The figure also shows two harmonics occurring at 60 and 90 Hz. It is this harmonic set of track laying tones that is detected by the ground array and is used to detect and follow the path of the vehicle. Since the exhaust signature of the vehicle contains less radiated energy than the radiated track/hull signature, and is facing away from the array in this run, the engine harmonics are not as intense. The fundamental engine firing frequency of 75 Hz and its harmonic set containing

150 and 225 Hz can be seen in the figure. Thus, both primary sources of vehicle signature, track and engine exhaust, need to be reduced to lower the detection of the vehicle. However, it is the track tonals that play the major role beyond 2000 feet, where atmospheric absorption will have reduced the engine harmonics and tracking becomes solely a detection of the track signature

THE BAND TRACK

The flexibility of the suspension wheels and the track shoes determine the magnitude of the impact forces for a given vehicle speed. The regular M113 track has inner compliant track pads that control the compliance between the suspension wheels and the track. The tested band track should show significant signature reduction because of its increased compliance and flexibility. However, this is offset by the inclusion of cleats in the band track which are required to provide the necessary traction. Thus, chordal action is not eliminated since the cleats themselves form a rigid series of "track shoes", which also generate chordal action. The purpose of the band track field test was to measure the effect of the reduced chordal action on the vehicle's signature as a function of vehicle speed. We can use equation 1 to determine the track laying frequency of the band track, since the M113 band track's cleat distance is 4.5inches. Thus, the band track fundamental track laying frequency

Frequency (Hz) = 3.9* V (mph)

Equation (3)

Figure 2 shows a spectrogram of the band tracked vehicle moving at 10 mph and detected at a distance of 500 feet

from the array. The harmonic set turns out to be 40, 80 and 120 Hz, very close to those predicted by equation 3.

THE ACOUSTIC GROUND ARRAY

The TARDEC acoustic ground array was used to detect the position of the vehicle and provide 3 second updates as the vehicle moved along the specified test path. Figure 3 shows the geometry of the 8 microphone acoustic ground array. The software for the array functions essentially as a delay sum beamformer, with delays introduced at each channel to bring the received signals in phase before summation. The array detected the harmonic line structure of the track and of the exhaust. There are two distinct areas of acoustic signature for the M113 vehicle and the demarcation is the 500 Hz line. Sources below 500 Hz make up the primary signature measured for ground vehicles, namely the track and exhaust set of harmonic lines that are measured by the ground array. Signatures of ground vehicles above 500 Hz, are due to the cooling fan and the whine of the transmission and are outside the frequency range measured by the array. Bearings of detection were measured with three second position updates and plotted as the vehicle moved along its driven path. The vehicle's motion was then determined relative to the sensor array's axis for several speed runs. The array was then moved to a different location to provide detection versus range data.

Since we stayed on the road behind the acoustic van, and the road curved away from a perpendicular to the vehicle path at stake #10, the detected path of the vehicle is curved because the line of bearing was closer to the array at the end of the vehicle's path when it was headed south (i.e. beyond stake #12). Likewise, the path was again curved when the vehicle started its northern run, since the vehicle was closer to the array at the beginning stages when it was headed towards stake #8.

The array's detection performance is shown in Figure 4 for a vehicle speed of 20 mph and a range of 1200 feet. The detection points were summed to achieve a numerical detection percentage. In the detection summation, the two detection points at 325 degrees and those that deviated from the vehicle's path by greater than 25 %, were not counted.

The array was also able to detect and follow other sources simultaneously during the vehicle's measurement. Thus airplanes that were passing overhead and the movement of wheeled vehicles in the area, were followed. Figure 5 shows the simultaneous detection of the standard M113 traveling south and the detection of a CUCV that was traveling to pick up personnel at the compound. It really didn't matter which direction the second vehicle was coming from, as long as the set of engine harmonics were within 4000 feet of the array.

The band track vehicle's signature was tracked by the existence of track cleats, which provided mobility and supplied the sources of chordal action. So, even though the band track signature was reduced, it was not eliminated. Figure 6 shows the detection of the band tracked vehicle moving at 20 mph and measured at 900 feet. Again a percentage of detection was determined for this speed and range.

RESULTS

The data collected by the array was arranged so that the percentage of

detection could be calculated for each vehicle speed and then comparisons were made between the regular and banded track. Figure 7 shows the detection versus range results for the two types of track for 10 mph. The vehicle exhaust is facing the array in this case. The data shows that the band track is 25 % quieter than the regular track. Figure 8 are the detection versus range results for the regular and the banded track for 20 mph with the exhaust facing the array. The band track is quieter by 35 %. The standard track signature is significantly greater at these higher speeds and this is also where the band trackshows important signature reduction.

CONCLUSIONS

The results of this study lead to the following conclusions:

1. The acoustic ground array detects the two primary sources of ground vehicle signature- the harmonic line structure of the track and the engine exhaust.
2. The track tonals are the dominant signature for vehicle speeds above 15 mph. For distances greater than 2000 feet, the engine tonals have been reduced by the atmosphere.
3. The signature reduction offered by the band track increases with vehicle speed, with 25 % reduction at 10 mph and 35 % at 20 mph.
4. The track tonals, due to their low frequencies below 100 Hz, will propagate across hilly terrrain and make the vehicle vulnerable at greater distances.
5. The track tonals are not eliminated by the band track due to the existence of cleats. The cleats are a

limiting factor for band track signature reduction.

REFERENCES

1. Prado G., " Acoustic Sensor Software User's Manual",May 1. 1996.
2. Shalis E, Wu Zhen, "Structural Acoustics Track Signature Model for Ground Vehicles: An Initial Study", 1995 GTM&V Proceedings, Houghton, MI
3. Galaitsis, A., "Computer Simulation Model for Double-Pin Tracked Vehicles", TACOM Tech Report #13117, July 1985.

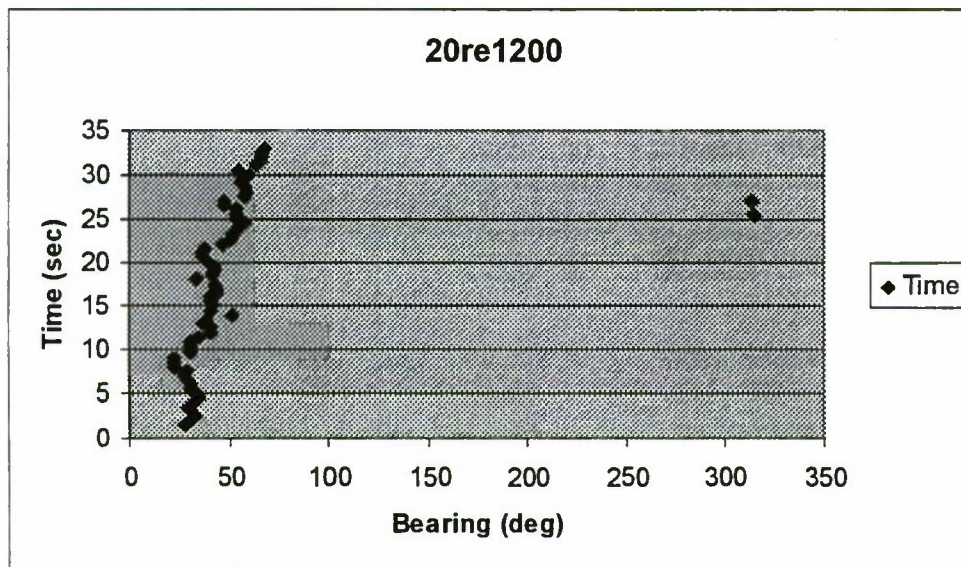**Figure 3. The geometrical ground array pointing north**

Figure 4. The regular M113 being tracked at 1200 feet for 20 mph.

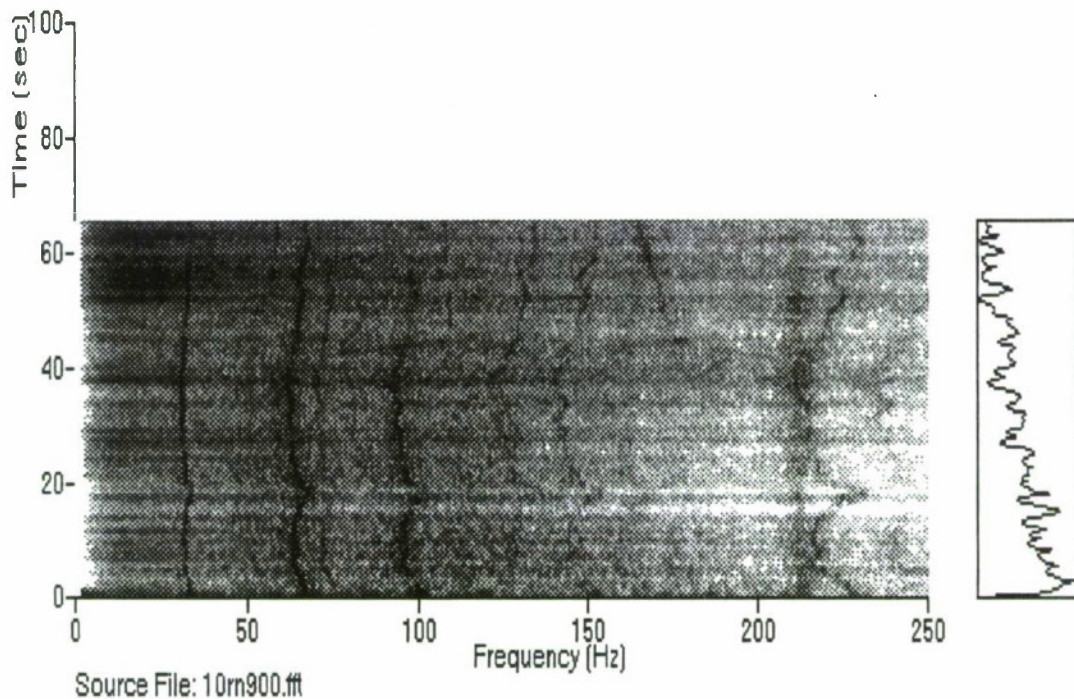**Figure 1** A spectrum of the baseline M113 measured at 900 ft while moving at 10 mph. The track tonal harmonic set is at 30,60 and 90 Hz. The engine fundamental firing frequency of 75 Hz and harmonics at 150 Hz and 225 Hz are used to follow the vehicle
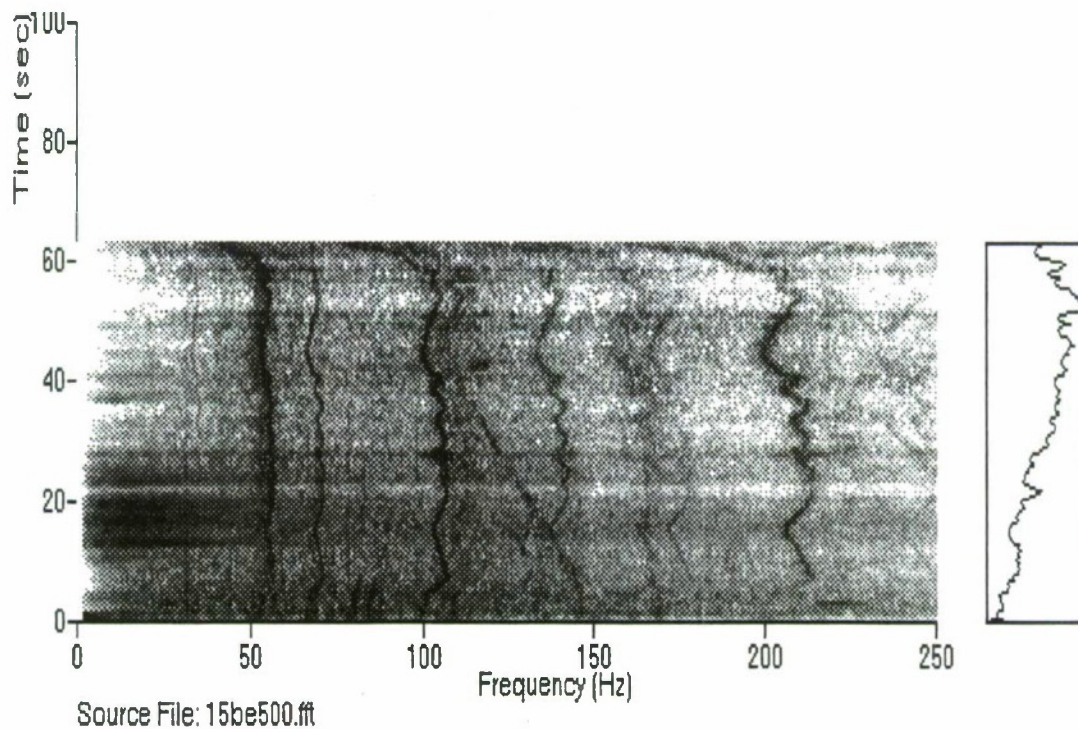
Figure 2 Spectrogram of the band tracked vehicle moving at 15 mph and detected by the ground array at 500 feet with the exhaust facing the array.

165

20re900
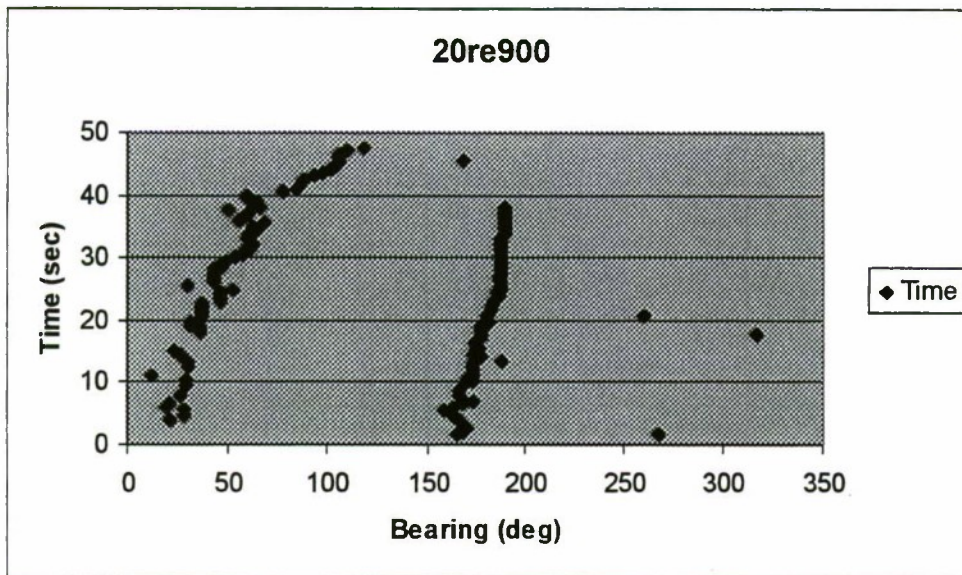
Figure 5. The ground array tracking two objects: the one on the left is the regular M113 moving at 20 mph.

20bn900

Figure 6. The band track vehicle detected at 900 feet

Regular M113 vs Band Track for 10mph (Exh Side)

Figure 7. Detection versus range comparison for the two types of track at 10 mph

Regular M113 vs Band Track for 20 mph (Exh Side)

Figure 8. Detection versus range comparison for 20 mph

Hyperspectral Algorithm Development for Army Applications

Marcos Sola and Dan Beekman
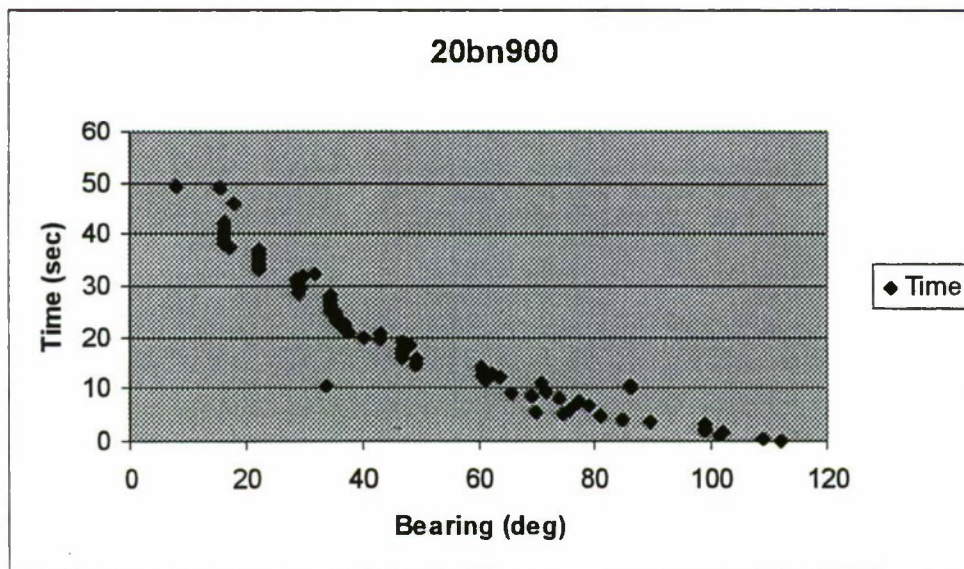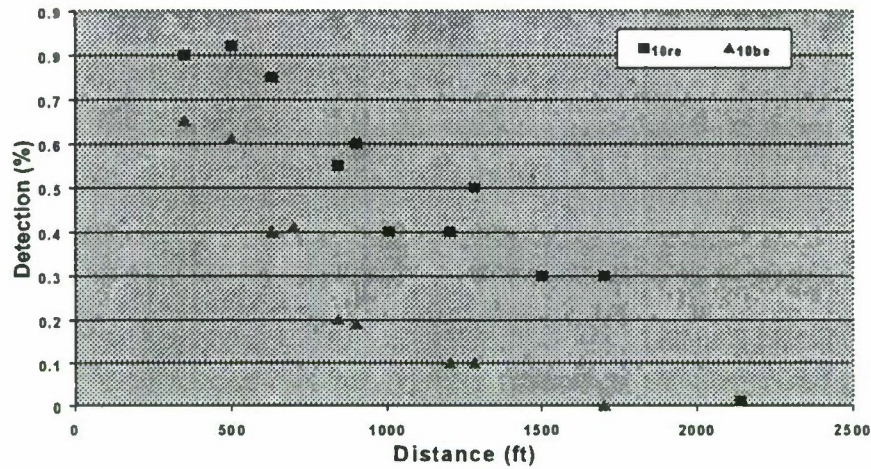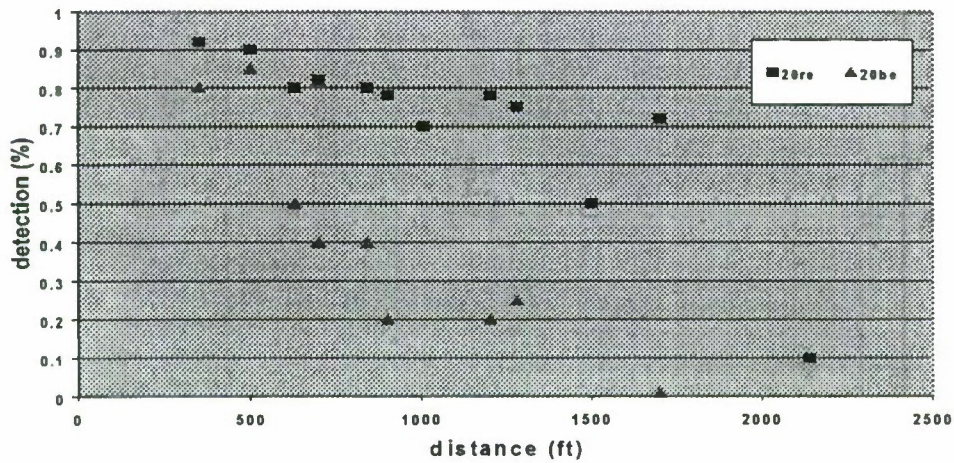U.S. Army Research Laboratory
Adelphi, MD 20793-1197

Abstract

We are adapting classic hyperspectral algorithm approaches to Army targeting applications. Principal component analysis and spectral matching are being combined with polarization and spatial information for the detection of targets in cluttered backgrounds. The ground-to-ground perspective introduces differences in targets and backgrounds compared with the usual high-altitude perspective of most remote sensing applications and databases. The acousto-optic tunable filter (AOTF) approach we have chosen also enables the use of multipass adaptive analysis in which a few spectro-polarimetric bands can be used for identifying regions of interest. Depending on the initial analysis, additional bands can be dynamically chosen to provide more discrimination. We acquired our data set with an AOTF sensor developed and tested by the Army Research Laboratory and the Carnegie Mellon Research Institute. The sensor provides imagery with variable polarization and spectral filtering over the visible and near infrared bands. Preliminary results obtained with a software image-processing system called Environment for Visualizing Images (ENVI) will be presented, along with plans for future collaborative activities within the Army Center of Excellence for Spectral Sensing Technology.

1. Introduction

Hyperspectral data have, for the most part, resided in the domain of the remote sensing community. These data, mostly in the visible through the short wave infrared and some in the mid- to far-infrared have been exploited for dual-use applications. In the commercial arena, these hyperspectral data (at the appropriate wavelength) have provided agricultural, geological, and terrain feature characterizations. For military applications, the hyperspectral data can detect (and possibly identify) chemical/biological agents and gaseous effluence from tactical targets, and serve as a countermeasure technology against high-value ground targets that employ camouflage, concealment, and deception (CCD) techniques. Figure 1 depicts the atmospheric
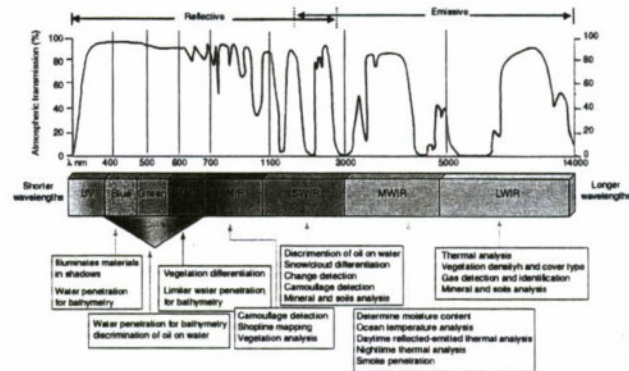

Figure 1. Atmospheric transmission as a function of wavelength and related applications.

transmission as a function of wavelength and shows the various regions of the electromagnetic spectrum and their intended application [1].

2. Army Requirements

For Army tactical applications, the interest in the sensor-to-target perspective is from ground-to-ground and from a shallow look angle, such as for a helicopter performing target acquisition at treetop level. We believe that there are differences in the requirements between what the Army needs versus the needs of the remote sensing community. Most of the data collection by the remote sensing community is done from a high altitude, looking down. The differences in these needs are shown in table 1; most noticeable is the small number of databases for ground-to-ground target acquisition data—this is a significant signature gap.

A top-down view of a target such as a tank does not need to convey much information for target acquisition. From a ground perspective, however, a target can present itself to the sensor in many aspect angles. Thus, a signature database would be quite large, and the computational time for signal processing may be significant for an automatic target recognition (ATR) sensor. Hyperspectral data collection by the remote sensing community is done via air breathers and

Table 1. Comparison between remote sensing community and Army tactical hyperspectral requirements.

Requirements	Remote sensing community	Army tactical
Angle of view	Top-down angle of view from an aerial platform	Ground-to-ground; shallow angle of view
ATR requirements	1 aspect angle	Multiple aspect angles (front, sides, and rear)
Primary function	Terrain categorization	Target detection, classification, recognition, identification
Acquisition level	Subpixel resolution	Several pixels on target
Background/clutter	Remote	Close-in
Atmospherics	Clouds, varying pressure level	Turbulence, dust, smoke, other obscurants
Signature databases	Many	Few-a significant data gap

(a)

(b)

Figure 2. Tanks under (a) test and (b) battlefield conditions.

satellites. The primary intent is to categorize terrain, vegetation, roads, or mineral deposits. This can entail subpixel classification depending on the altitude and resolution of the sensor. The Topographical Engineering Center (TEC), an Army agency, falls under this category. The decision to fire a missile on a target will not be made based on spectral information alone unless the data have been made part of a data fusion schema. In a tactical engagement, the hyperspectral imager (HSI) will have many pixels on target. Hence, the HSI provides the shooter with both *spectral and spatial* features of the target signature. This is a desirable feature for counter-counter measure (CCM) against tactical targets operating in a realistic battlefield scenario (see fig. 2).

Besides the difference in the background and clutter due to the line-of-sight perspective, we believe that the impact of atmospherics on target acquisition will also be different. For ground-to-ground target acquisition, turbulence could prove to be more of a problem than for remote aerial sensing. There is a hyperspectral signature database gap for Army tactical requirements. This gap, especially in the mid- to the long-infrared wavelengths needs to be populated with high-quality signature measurements to be usable for HSI design, prototype enhancements, and algorithms and modeling and simulation development as well as for phenomenological understanding of the signature dynamics. These research and development initiatives do impact measurement and signature intelligence (MASINT) capabilities in the near term.

3. Army Center of Excellence

Within the Department of Defense, the Deputy Director for Research and Engineering (DDR&E) has created an Integrated Process Team (IPT) to exploit the hyperspectral data for Army applications. In 1997, the IPT recommended that the following key elements to be addressed:

- Target signature libraries: To establish a library of key target and feature signatures.

- Background signature libraries: To build a library of well-characterized backgrounds and environments.

- Algorithm development and evaluation: To develop and evaluate algorithms for processing hyperspectral data and target detection.

- Sensing performance limits: To determine attainable HSI performance.

- Predictive model validation: To define necessary and sufficient metrics for a robust model validation.

In response to these recommendations, an Army Center of Excellence was established [2] that consists of TEC,

169

the Night Vision and Electro-Optics Directorate (NVESD), the Space Missile Defense Technology Center (SMDTC), and the Army Research Laboratory (ARL). Within ARL, the participants are the Army Research Office (ARO) and the Sensors and Electron Devices Directorate (SEDD).

4. ARL's Hyperspectral-Related Projects

Our interest and investment in the hyperspectral arena complements the findings of the IPT. We are performing HSI prototype development and characterization using acousto-optical tunable filter (AOTF) technology. We have performed field experiments on the visible to near-infrared and short wavelength infrared (SWIR) in conjunction with a visible to near-infrared grating spectrometer from the Naval Research Laboratory (NRL) to gather data. The intended use of these data is to further HSI enhancements and signature analysis to support signature phenomenology understanding, modeling and simulation, and ATR algorithm development and to ease the computational burden in signal processing for target acquisition. We are also developing dual-color broadband FLIRs via our Federated Laboratory program—FedLab [3]. Our goal is to be able to identify hyperspectral technologies for Army tactical applications and identify data gaps.

4.1. Comparison of Hyperspectral Systems

Hyperspectral imagers can be based on a variety of spectral dispersing elements; for example, gratings, prisms, Fourier transform infrared (FTIR) spectrometers, or acousto-optic tunable filters (AOTFs). Each of these systems is subject to constraints, such as small angular acceptance (large f number) and narrow spectral passband, that limit the photon flux and decrease the detector signal-to-noise ratio (SNR) compared to the broadband case. Some SNRs can be regained by increasing the integration time or cooling the detector.

The systems differ in the method that they use to create an x-y-λ data cube with a two-dimensional detector array. Both the grating and prism systems acquire y-λ images and scan the x-dimension with scan mirrors or platform motion to create the x-y-λ cube. The AOTF system acquires x-y images and scans λ in successive frames. The FTIR system acquires x-y interference patterns and scans the spectrometer mirror to create an interferogram for each pixel. An inverse Fourier transform is performed to create the x-y-λ data cube. Each of these techniques has benefits and drawbacks, and the choice of system depends on the application.

5. Hyperspectral Signature Analysis

We are working with our FedLab partners but taking a different approach in the signature analysis of a hyperspectral signature (HSS) database. Our goal is to identify for each class type and class set the specific bands for which there is a unique feature in spectral response. We then intend to use band ratios or image subtraction plus other nonintensive computational algorithms to extract the target from the background/clutter. We want to be able to categorize unique hyperspectral bands over a wide dynamic range of the HSS database suited for Army tactical requirements. This approach can use the advantage that AOTF technology offers for selecting specific bands for target acquisition signal processing in real time.

5.1. Image cube

A software image-processing system called Environment for Visualizing Images (ENVI) [4] was used to generate the image cube shown in figure 3. These data were taken with a visible to near-infrared AOTF sensor at Fort AP Hill, VA [5]. For a given pixel (x,y) in the scene, we follow the usual convention of the three-profile plot definition. The x- and y-profile distribution provides the reflectance as a function of sample and detector line as shown in figure 3. The plot is for a single hyperspectral wavelength. Figure 3 also shows the z-profile, which gives the reflectance for a single selected pixel as a function of the hyperspectral wavelength that is layered in the z direction.

5.2. AP Hill Data (June 1998)

Figure 4 (a) is an image scene at 760-nm wavelength. Several regions of interest (ROIs) were selected and plotted as shown in figure 4 (b). The number of sampled pixels
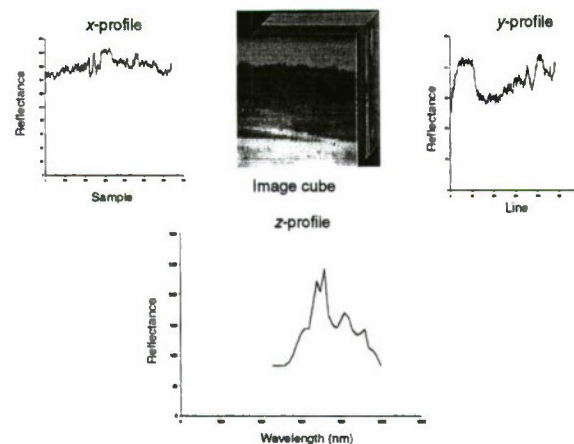

Figure 3. Image cube.

for each different ROI was not kept constant but the number was at least greater than 50. Other than the difference in the reflectance value for some ROIs, the shape of the curves appeared to be similar. This is probably due to the large number of pixels sampled for a given ROI. In doing this, the resolution of the purity of the pixels chosen to belong to a given class was degraded. If we take the peak of the reflectance response at λ_{760}, subtract the next peak from it, and then apply a low-pass filter, we obtain figure 4 (c). In this figure, as compared to figure 4 (a), there appears to be an area of interest in the search sector to which other fused sensors can be directed for further data interrogation. Figure 4 (d) is a camcorder image of a high-mobility multipurpose wheeled vehicle (HMMWV) at a range of 1.2 km. The vehicle is presented at its right-side aspect angle. A similar figure can be inferred from the processed image (see fig. 4 (c)). This technique is appealing because the computational burden to process the data was not intensive. However, performing these operations for other peaks, other polarization angles, and other targets at different scenarios did not provide any interesting information. A single pixel sampling for the various parts of the HMMWV target is shown in figure 5. To enhance the analysis for determining if there are differences in the shape of the curves, the plots were stacked as shown in figure 6 so that the y-axis value is offset for clarity. In the spectral band between 550 and 650 nm, a difference in the slope of the curve between the front and rear tires versus the rear panel and the Plexiglas window can be inferred. This is an example of a spectral relationship between different classes that can be categorized and studied to see if spatial information can be obtained.
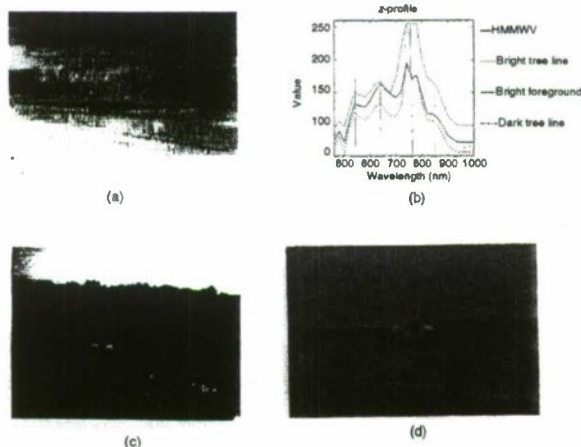
Figure 4. (a) Image scene of an HMMWV at 760 nm wavelength, (b) plots of several regions of interest, (c) image scene after applying low-pass filter, and (d) camcorder image of HMMWV.
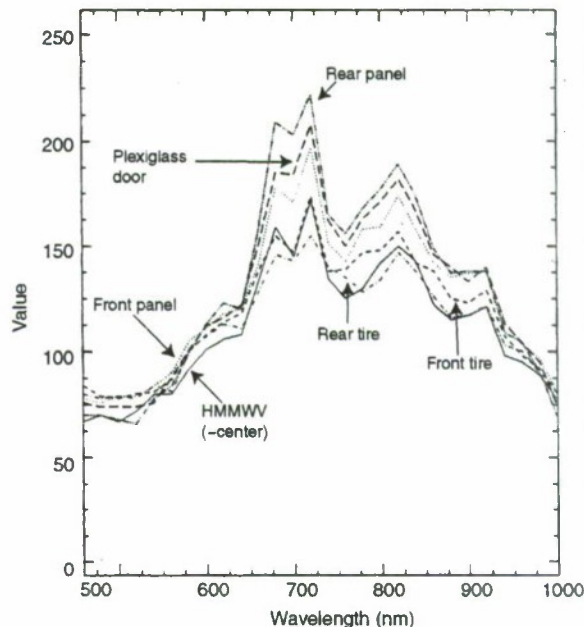
Figure 5. Single pixel sampling for parts of HMMWV target.

Figure 6. Stacked plots of various parts of HMMWV target.

Figure 7 shows a principal component analysis (PCA) performed on the data representing 45° polarization measurements. Most of the information for the scene is con-

171

tained within the first three bands. Figure 8 (a) is the original scene at λ_{760} as compared to figure 8 (b), which is the first principle component; figure 8 (c), the fifth principle component; and figure 8 (d), the sixth principle component, where the scene is mostly noise. If the target's signature falls along the noise level of the backgrou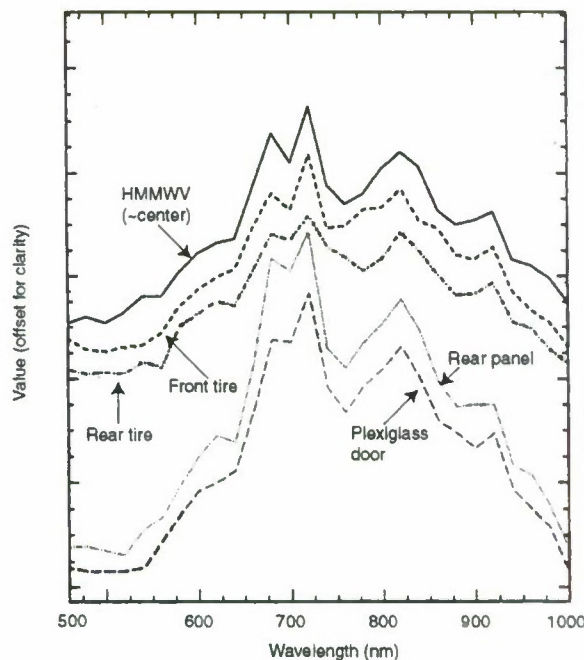nd/clutter (as with a low observable or a target using CCD), detection using this technique alone will provide a high false-alarm rate.

Figures 9 (a) through (d) are image scenes at 760-nm wavelength and polarization angles of 0°, 45°, 90°, and 135°, respectively. A z-profile taken at about the centroid of the HMMWV is shown in figure 10. The z-profiles for 45°, 90°, and 135° polarization angles are similar. This group, on average, indicates some differences to that for 0° polarization. Using this spectral information at λ_{760}, a low-pass filter was applied to the image difference between 45° and 0° polarization. This yielded the spatial information shown in figure 11. The scene shows a region of interest (side aspect of a vehicle) that other sensors could be vectored into to gain more data on the suspected target.

Figure 7. Principle component analysis of data taken at 45° polarization angle.

Figure 9. Image scenes at 760 nm wavelength and polarization angles of (a) 0°, (b) 45°, (c) 90°, and (d) 135°.

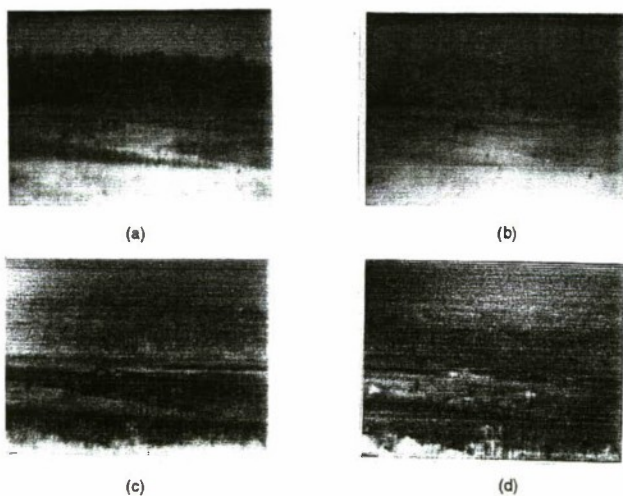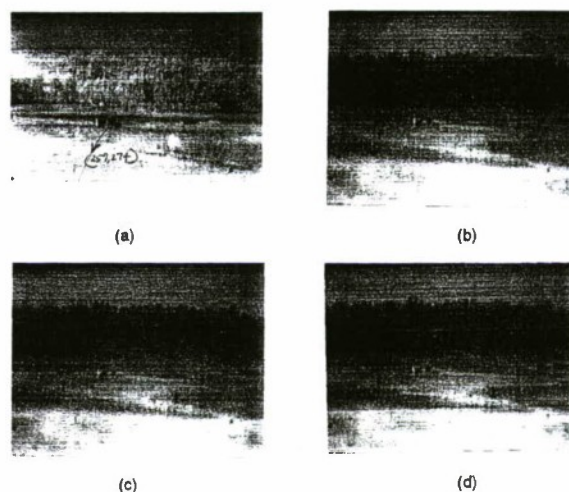Figure 8. (a) Original scene taken at Fort AP Hill, VA, and (b) first principle component, (c) fifth principle component, and (d) sixth principle component.
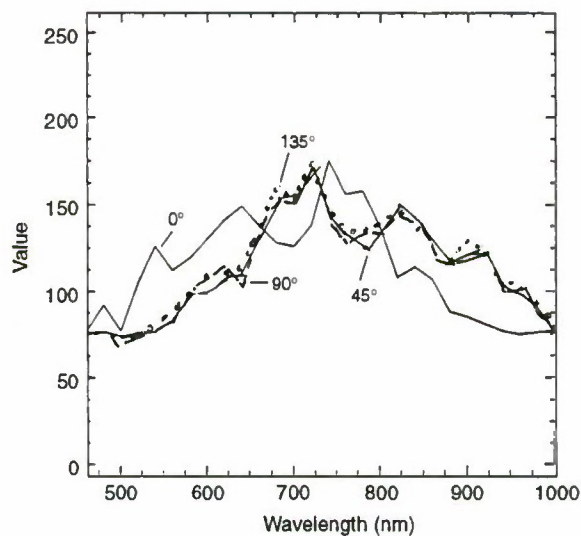

Figure 10. Z-profile at about the centroid of HMMWV.

172
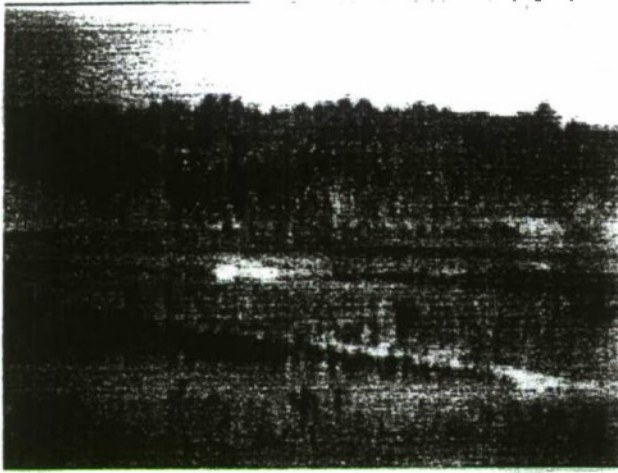
Figure 11. Image subtraction at λ_{760} and low-pass filter with (45° - 0°) polarization.

Figure 12. Z-profile at various single pixel regions of interest taken from an LWIR sensor image.

5.3. LWIR Data Analysis

Figure 12 is a z-profile at various single pixel ROIs taken from an image by an LWIR sensor [6]. Other than the obvious differences in the spectral radiance generated using a single pixel for different ROIs, the curves look similar. We need to be able to plot the z-profile with a y-axis that shows the most sensitivity as a function of wavelength for different ROIs. The approach taken is to look at emissivity as a function of wavelength using the measured spectral radiance. Consider the measured spectral radiance $L_M(\lambda)$. For a given snapshot in time, range, and atmospheric transmission, we represent the measured spectral radiance as

$$L_M(\lambda) = L_{BB}(\lambda)\,\varepsilon\,(\lambda)X_{Mis}, \qquad (1)$$

where

$L_{BB}(\lambda)$ = blackbody curve,

$\varepsilon(\lambda)$ = emissivity, and

X_{Mis} = atmospheric transmission,

  but

$$L_{BB}(\lambda) = 2c^2h/\lambda^5(e^{ch/\lambda kT}-1)\text{(Planck's equation),} \quad (2)$$

where

c = 2.99793E + 08 m/s, speed of light,

h = 6.62620E − 34 J/s, Planck's constant,

k = 1.308062E − 23 J/K, Boltzmann's constant, and

T = 273 + °C, absolute temperature in degrees Kelvin (K).

From the measured data, find the wavelength (λ_{Max}) that provides the maximum spectral radiance. Use this wavelength in Wien's displacement law:

$$\lambda_{Max}T = 2898 \; \mu\text{m K}. \qquad (3)$$

T can be solved and used in equation (2) to generate the blackbody curve. If the blackbody radiation is taken at the same range (as the target for example) and degraded by the same atmospheric transmission, the apparent blackbody equation ($L_{ABB}(\lambda)$) is given by

$$L_{ABB}(\lambda) = X_{Mis}L_{BB}(\lambda). \qquad (4)$$

Dividing equation (1) by equation (4) yields

$$\varepsilon(\lambda) = L_M(\lambda)/L_{ABB}(\lambda). \qquad (5)$$

We can now generate the emissivity as a function of wavelength. For the data analyzed, the slant path from the tower (approximately 100 m high) to the open target was approximately 500 m and was taken in the summer. For the purpose of this analysis, we set the atmospheric transmission to be unity. (The spreadsheet allows for any input in atmospheric transmission as a function of the wavelength if these data were collected as part of the ground-truth measurements.) A single pixel spectral radiance was generated for one target and three different types of background (a dirt road, a shaded tree, and an area shadowed by trees). The digitized measured data showed a maximum spectral radiance along with the calculated absolute temperature and are shown in table 2.

Figure 13 shows preliminary data using this emissivity approach. Relative emissivity is plotted as a function of wavelength. The shape of each curve is more important (at least for now) than the absolute value. Our next task will be to identify bands of interest and attempt to extract spatial information to locate the target.

173

Table 2. Four object classes were sampled using a single pixel to generate the z-profile. The absolute temperature was from Wien's displacement law using the hyperspectral wavelength that produced the maximum measured spectral radiance.

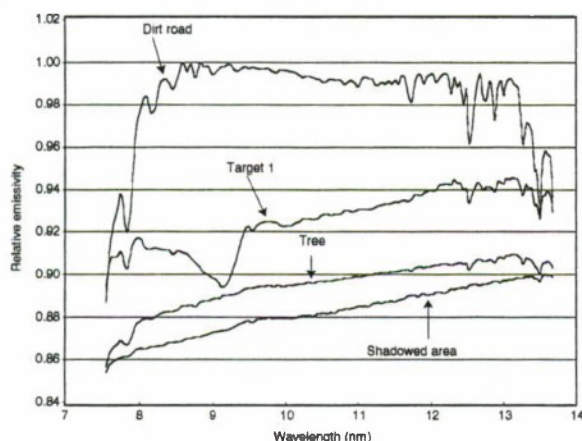| ROI	λ_{Max}	Band No.	$T(K)$
Dirt road (72, 73)	9.7638	38	296.81
Tree (72, 41)	9.7638	38	296.81
Area shadowed by trees (72, 56)	9.7638	38	296.81
Target 1 (72, 65)	9.8165	39	295.22

Figure 13. Preliminary data using emissivity approach.

5.4. Other Algorithms/Tools

There are numerous hyperspectral algorithms/tools that are available. Both Applied Signal and Image Technology (ASIT), Inc., [7] and NRL make use of an anomalous detection algorithm. This is further enhanced using algorithms developed internally within the respective company or laboratory. ASIT's software and NRL's Orassis software showed good results on the respective database that the software was tested against. Figure 14 shows ASIT's algorithm identifying targets/features of interest. The data were taken with a visible to near-infrared sensor from an overhead platform. The targets/features identified include U.S. military vehicle paint under various conditions: by itself, under camouflage, and under camouflage/vinyl nets. Also identified are beach defenses and barricades as well as some spectral anomalies. Both algorithms were not able to find the target in the AP Hill experiment that was described earlier in this paper. Whether this is a problem with our data set, a glitch in the software used, or simply because the AP Hill data

showed a target that is part of the principal components that an anomalous detection algorithm will reject is not known. A generic problem with most algorithms is that they usually perform well on the database that they are trained on. However, over a wide dynamic signature range, the algorithm's performance has the potential of deteriorating to an unacceptable level. Another algorithm developer that we are working with under a FedLab agreement is ERIM. ERIM is taking several approaches. These include neural networks, a pairwise adaptive linear matched filter, and apparent temperature versus wavelength. In addition, TEC has large holdings in a hyperspectral database library that includes calibrated signature data on various trees, grassy roads, materials, and other terrain features.

6. Conclusions

The following preliminary conclusions have emerged based on our work to date:

- There does not appear to be an algorithm that would detect/identify tactical targets over a wide dynamic signature range. Most of the success is from a database that the algorithm was trained for.

- For Army tactical applications, there is a hyperspectral signature database gap for the ground-to-ground perspective in the visible through the long wavelength infrared spectral band.

- Hyperspectral imagers can provide spectral and spatial information to enhance target discrimination.

- AOTF imagers are rugged fieldable systems for obtaining hyperspectral data.
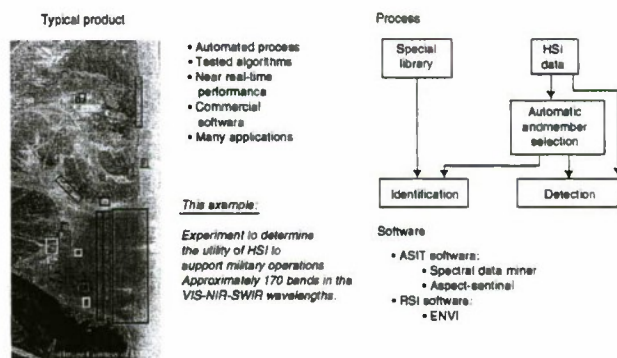
Figure 14. ASIT algorithm results of identified targets/features of interest.

174

References

[1] MASINT Spectral Processing Course, Spectral Information Technology Applications Center (SITAC), Fairfax, VA (2–6 November 1998).

[2] Hamilton, M., and Roper, W., "U.S. Army's Center of Excellence for Spectral Sensing Technology," F. Sadjadi, ed., *Proceedings of the SPIE on Automatic Target Recognition IX*, **3718** (1999).

[3] Pollehn, H. K., and Ahearn, J., "Multi-Domain Smart Sensors," B. Andresen and M. Scholl, eds., *Proceedings of the SPIE on Infrared Technology and Applications XXV*, bf 3698 (1999), pp 420–423.

[4] Environment for Visualizing Images (ENVI), Version 3.1 software, Research Systems, Boulder, CO.

[5] Gupta, N., *Proceedings of the First Army Research Laboratory Acousto-Optic Tunable Filter Workshop*, ARL-SR-54, Army Research Laboratory, Adelphi, MD (March 1997).

[6] Brassboard Airborne Multispectral Sensor System Specifications Program (BAMS4), data collection from the Robert F. Russell Measurement Facility, Redstone Arsenal (9–17 October 1996).

[7] Applied Signal and Image Technology, Inc., http://www.asitinc.com/ (August 1999).

Imaging Spectroscopy for Signature Modeling

J.P. Garcia, C.E. Volin, M.R. Descour, E.L. Dereniak, D.S. Sabatke
Optical Sciences Center
University of Arizona
Tucson, AZ, 85721

David T. Sass
U.S. Army TACOM
Warren, MI 48397

Chris G. Simi
NVEOL
Ft. Belvoir, Va. 22315

ABSTRACT

Imaging spectroscopy with a simultaneous spectral/spatial infrared sensor is described and discussed. This new technique uses spatial information on a two dimensional focal plane to calculate the spectral content of the image. The object cubes produced are available at 30 fps. A MWIR instrument will be described which could be used to identify fuel mixtures and exploding ordinances.

INTRODUCTION

Collection of spatial and spectral information has both military and commercial applications, where not only the location of the event is important to the observer but also the spectral content for identification. Up to now, simultaneous imaging and spectrometry were impossible. Spectral images were obtained by scanning in time or space. Recent research developments have produced an imaging spectrometer, which simultaneously collects both spectral and 2D spatial information about a scene[1]. To demonstrate our concept, it helps to think of a hyperspectral object cube as shown in Fig. 1.
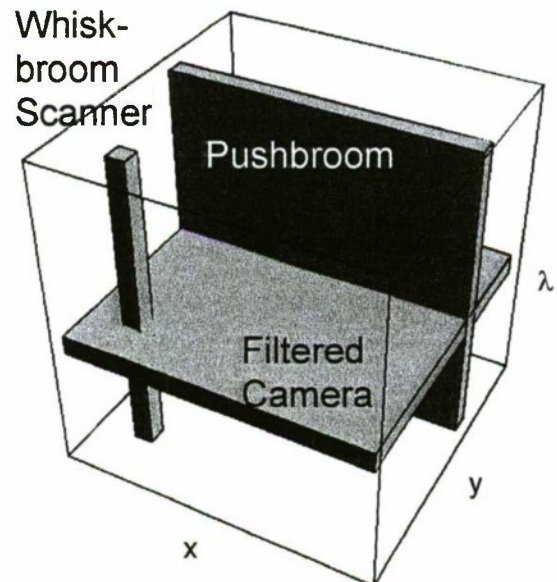

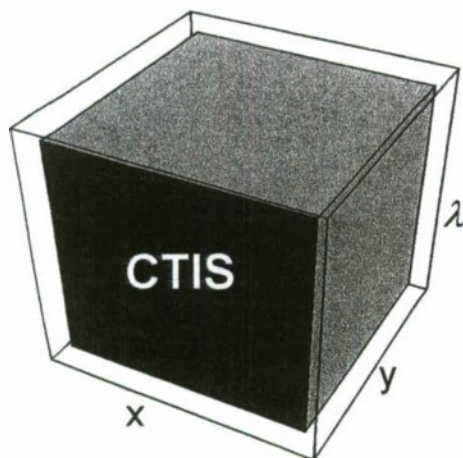Fig. 1. Object cube collected conventionally.

Fig. 2. Reconstructed CTIS snapshot object.

Fig. 3. Projections of the object cube onto the focal plane. Spectral information along the vertical axis of the object cube is projected along the radial coordinate of the focal plane and multiplexed with spatial information.

The object cube represents the x, y, λ information of a scene. Any slice through the cube parallel to the x-y plane represents a monochromatic image. Shown in Fig. 1 are different conventional approaches to acquiring the data within the object cube. Note that they all require SCANNING to fill the three dimensional space. The Computed Tomographic Imaging Spectrometer (CTIS) obtains information about the entire object cube each frame time as shown in Fig. 2. This allows reconstruction of spatial and spectral information for the scene at each frame time, which is essential for scenes that change in space/wavelength with time. The CTIS approach to hyperspectral imaging was facilitated by the development of large two-dimensional arrays and fast computing capabilities. In order to acquire the spectra of the scene on a broadband detector array, an imaging technique, based on computer axial tomography (CAT), was adopted. This combination of technologies allowed the development of a no moving parts, optically simple instrument called CTIS[2].

OVERVIEW OF CTIS INSTRUMENT

The driving concept behind the development of the CTIS is the reconstruction of a 3D object from 2D projections. The objective here is to record 2D projections of the x, y, λ object cube. If sufficient projections are recorded, then the original object cube can be reconstructed. These projections are shown schematically in Fig. 3. A two-dimensional array records these many "shadows" and reconstruction of the (x, y, λ) object is accomplished via inversion methods.

The spectrometer consists of three optical-element groups: an objective lens, a collimator lens, and a re-imaging lens.
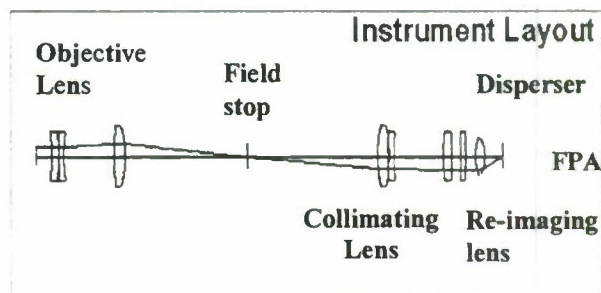
Fig. 4. Optical schematic of the CTIS

The use of AR coated lenses for the 3 − 5 μm band in the optical train greatly enhances the performance of the instrument. The Computer Generated Hologram (CGH) disperser[3] is located in collimated space between the collimator lens and the re-imaging lens. The FPA in the prototype system shown in Fig 5. is a160 x 120 InSb array with 50 μm pitch and 30% fill factor operating at 30 frames per second and 5 msec. integration time. The Stirling cycle cooler built in to the camera enhances portability of the system. The 5mm square field stop maps to a 20 x 20 pixel area on the FPA.
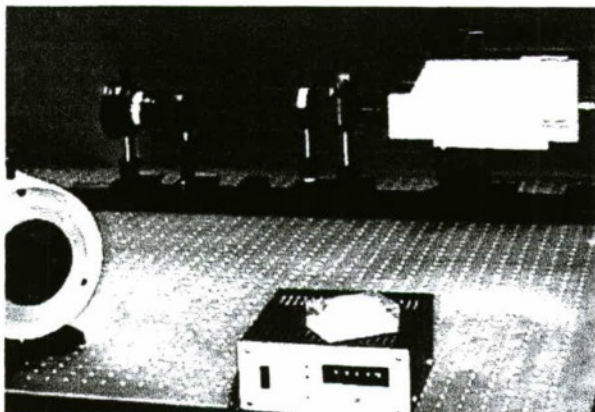
Fig 5. Prototype MWIR CTIS. In the background, from left to right, a 100 mm focal length objective lens focuses the target onto a 5mm square field stop. The field stop is appropriately minified and dispersed on to the InSb FPA by the combination of a 200 mm collimating lens, GaAs binary diffraction grating and 50 mm re-imaging lens.

While the IR lenses were purchased from vendors, the two-dimensional binary disperser, which is actually a diffractive phase grating, was custom fabricated at the Optical Sciences Center in GaAs. The phase grating shown in Fig. 6, which has a period of 90 µm in both the x and y directions, was chosen in combination with the field stop size, the collimating lens focal length and the re-imaging lens focal length to place the zero and first diffraction orders with in the FPA periphery.
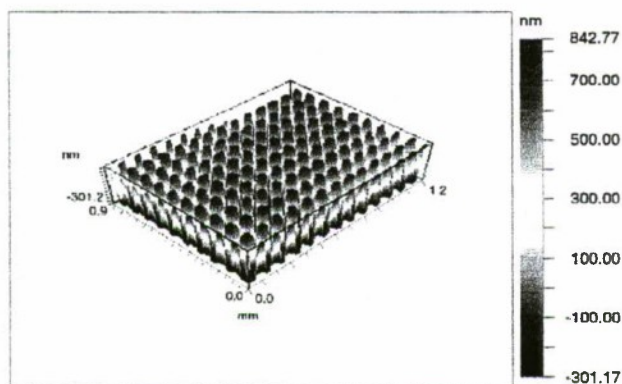
Fig. 6. Wyko interferometer surface profile of the 2-D binary phase grating fabricated using reactive Ion Etching at the Optical Sciences Center by Michael Descour and Daniel Simon.

The physical depth of the grating was initially chosen to be 0.7 µm based on a model derived from Goodman's treatment of a sinusoidal phase grating[4]. The depth was subsequently fine tuned by trial and error to 1.1 µm to give reasonable irradiance uniformity across he FPA for a 973 K blackbody target.

CALIBRATION

The MWIR CTIS calibration process is similar to the visible CTIS calibration process. The output port of a 973K blackbody is coupled directly into a computer-controlled monochromator with a 4-µm blazed grating. The monochromator allows selection of spectral bands as narrow as 0.1 µm. Since we do not have a suitable fiber light guide or coupling lenses available for such a waveguide, the output of the monochromator must be imaged into the field stop (via a fold mirror) from a distance using the spectrometer's objective lens. The exit slit of the monochromator is stopped down to a square with an opaque aperture.

A fold mirror is used to position the calibration spot in the field stop so that it exactly fills four (50-µm) pixels on the FPA. The monochromator is then scanned through wavelengths and a calibration image is acquired at each spectral band. Since the instrument is essentially shift invariant over each wavelength, the calibration images may be software-shifted to fill the field stop. There is currently no allowance for absorption over the path difference external to the objective, but absorption over such a short path should not be significant outside of the CO_2 absorption band from 4.2 to 4.4 µm. There is also no normalization of the calibration images by the irradiance of the incident light.

It is not yet established whether the objective lens is exchangeable, as in the visible instrument, under these calibration conditions. Assuming low aberrations from 2.95 to 4.55 µm, (especially distortion) in the objective lenses, it should be possible to exchange objectives without losing calibration accuracy. It is very simple to factor out the ratio of the spectral transmission of the exchanged objectives from the calibration images. This technique has been used to calibrate the MWIR CTIS with $10 \times 10 \times 8$ (x,y,λ) voxels. The spectral sampling is 2 physical pixels in size, which is the limiting spatial factor for this type of calibration.

RESULTS

Preliminary results have been obtained using laboratory targets such as blackbodies, blackbodies through narrowband filters and combustion products from a butane lighter.

a

b

Object: 550°C (823 K) blackbody at 2 meters

CO_2 absorption band (4.2 - 4.4 μm)

Blackbody

Lighter
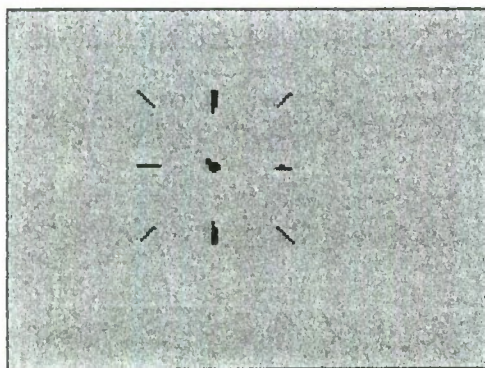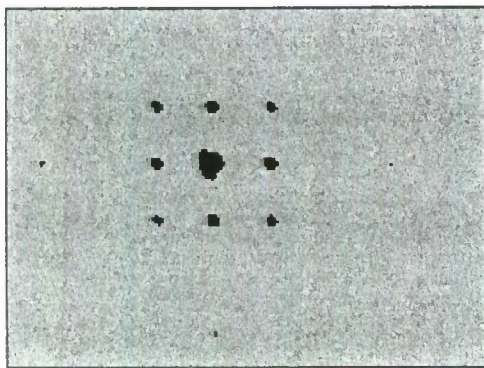
Fig. 7. Raw images of laboratory test targets: a) 350 C blackbody; b) 3.3 μm narrowband filter in front of a blackbody. The image contrast has been reversed; the distance from the center order axis is proportional to wavelength.
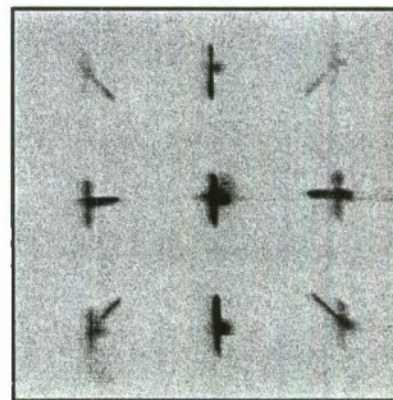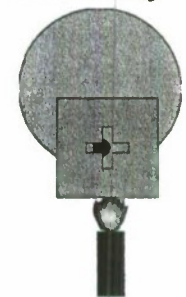
Fig. 8. Top: MWIR CTIS snapshot showing 823 K blackbody spectrum with CO_2 absorption band. Bottom: MWIR CTIS snapshot of an 823 K blackbody combined with hot CO_2 from a lighter seen through a cross-shaped aperture. The center of cross appears near the location of the CO_2 absorption band. A 256 x 256 PtSi camera was used in a CTIS configuration for these two images.

179

RECENT IMPROVEMENTS

A 299 mm objective lens and computer camera control have been recently implemented. By placing the camera's gain and offset settings under computer control, we have been able to greatly improve the system's sensitivity to objects heated less than 100 C above room temperature. Shown below is a photo of a black paint test target which consists of a piece of aluminum painted on the left side with Krylon ® Ultra Flat Black paint (#1602) and Krylon ® Semi Flat Black paint (#1613) on the right side. The left edge of the vertical black bar marks the dividing line between the two paints. The samples were heated to approximately 70 ^{0} C on the hot plate and the commercial blackbody was set for 200 ^{0} C.
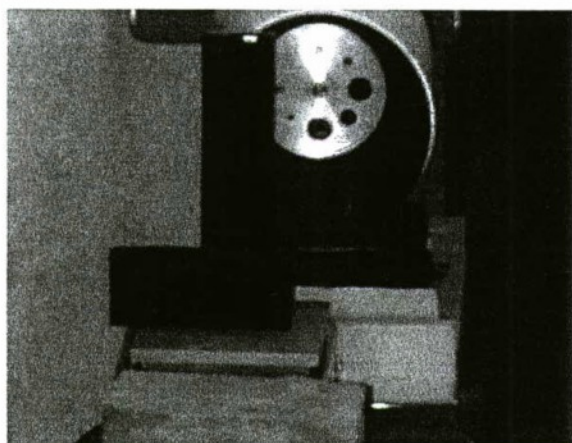
Fig. 9. The black paint test target is the horizontal bar on the hot plate. The left edge of the vertical black bar marks the dividing line between Krylon ® Ultra Flat Black on the left and Krylon ® Semi Flat Black paint on the right.

Fig. 10. MWIR CTIS raw focal plane image of the black paint test target. The four brightest projections are arranged about the center order on roughly 45 degree diagonals. The bright streak in each of the projections is the 200^0 C blackbody. Background subtraction has been performed.

The raw MWIR snapshot image of Fig. 10 demonstrates the spreading of the broadband image in the center into four primary diffracted orders which disperse the object's spectrum in the radial direction. Note that virtually no difference between the paints can be seen in the center image. The four brightest projections and the EM algorithm[5] were used to reconstruct the spectral-spatial data cube. One spectral slice is shown in Figure 11. Resels A and B were chosen as representative of the Ultra Flat Black paint side; resels C and D were chosen as representative of the Semi Flat Black paint. Both paint samples should be at similar temperatures since they were sprayed onto a common slab of aluminum. Resel E is the sub-resolution blackbody source.

Fig. 11. 3.95 – 4.15 µm spectral slice of the reconstructed object cube. Resels A and B are Ultra Flat; C and D are Semi Flat. Resel E is the blackbody.

In Figure 12, the relative spectra of resels A, B and E are plotted against the band center wavelength; in Figure 13, the relative spectra of resels C, D and E are plotted in the same fashion.
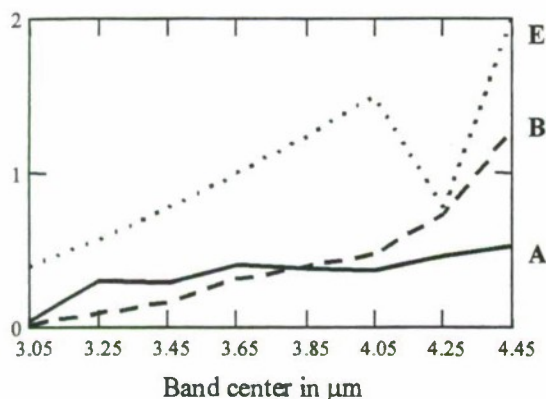
Fig. 12. Relative spectra of resels A and B on the Ultra Flat Black sample versus band center wavelength in microns. Resel E is the 200^0 C sub-resolution blackbody.
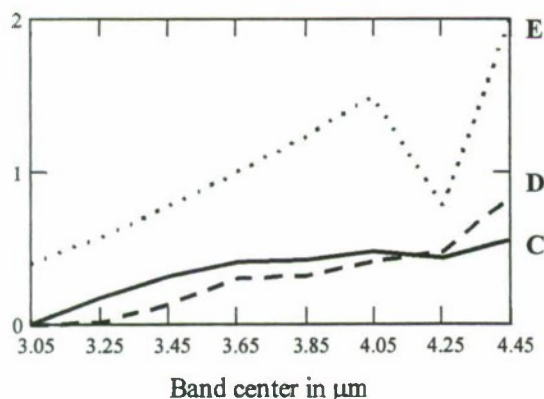
Fig. 13. Relative spectra of two resels C and D on the Semi Flat Black sample versus band center wavelength in microns. Resel E is the 200^0 C sub-resolution blackbody

Inspection of these two figures reveals subtle yet significant differences in the spectra of the two paint samples. In the 3.85 - 4.25 μm region, the Ultra Flat Black sample exhibits a shallow downward inflection; whereas, the Semi Flat Black sample shows a slight upward inflection. The Semi Flat Black sample has a distinct downward inflection centered in the 4.25 μm band; whereas, the other paint does not. Even though CO_2 absorption exerts an influence on the data, as evidenced by the dip of the blackbody curve in the 4.25 μm band, the influence is equal for all portions of the reconstructed image since the target range of 14 meters is essentially constant for all resels. Therefore the relative spectral features observed above are truly indicative of sample differences. Thus the instrument has the capability to distinguish between objects which look identical to a broadband MWIR imager. The magnitudes of the spectral curves in Fig. 12 and Fig. 13 are strongly dependent on the extent of the source compared to a resel IFOV of 0.8 inches in this experiment. The 0.200 inch exit aperture of the blackbody reduces the magnitude of its spectrum compared to other resels.
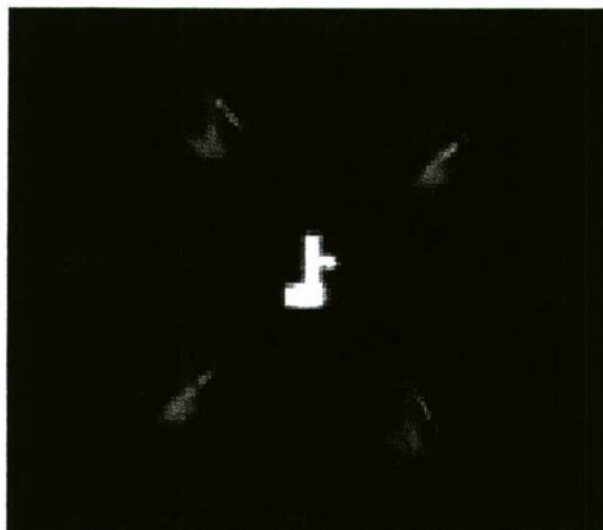
DISCUSSION

The instrument could be used to distinguish between friendly and foe vehicles at nearly 30 fps when the targets are too distant to be sufficiently resolved for shape identification. Reconstruction time for the entire 10 x 10 x 8 object cube is currently less than 30 milliseconds on a 333 MHz Pentium II machine and anticipated advances in PC processing power should make larger cubes available in near real time. With a 256 x 256 array and a well-designed CGH disperser, 20 x 20 x 15 reconstructions are practical.

The instrument has also been proposed for use in armor kill assessment tasks. The presence of burning fuel versus burning propellant is the distinguishing feature of a soft versus a hard kill. In the 3.5 – 4.2 micron region, the combustion products from these two sources show sufficient spectral variation to make the problem tractable with as few as five spectral bands[6].

CONCLUSION

We have demonstrated a high-speed snapshot imaging spectrometer prototype operating in the MWIR. Relative spectral mapping of an object has been demonstrated, although an absolute radiometric calibration of the instrument has not yet been performed. Future work includes performing this calibration to enable absolute temperature mapping and absolute spectral characterization.

ACKNOWLEDGEMENTS

REFERENCES

1. M.R. Descour, C.E. Volin, T.M. Gleeson, M.F. Hopkins, D.W. Wilson, P.D. Maker, "Demonstration of a Computed-Tomography Imaging Spectrometer using a computer-generated hologram disperser", *Applied Optics*, **36**, No. 16, pp. 3694 – 3698 (June 1, 1997).

2. M.R. Descour, C.E. Volin, E.L. Dereniak, K.J. Thome, A.B. Schumacher, D.W. Wilson, P.D. Maker, "Demonstration of a High Speed Non-scanning Imaging Spectrometer", *Optics Letters,* **22**, No. 16, pp. 1271 – 1273 (August 15, 1997).

3. M.R. Descour, C.E. Volin, T.M. Gleeson, M.F. Hopkins, D.W. Wilson, P.D. Maker, "Demonstration of a Computed-Tomography Imaging Spectrometer using a computer-generated hologram disperser", *Applied Optics*, **36**, No. 16, pp. 3694 – 3698 (June 1, 1997).

4. J.W. Goodman, Introduction To Fourier Optics, McGraw-Hill, San Francisco, 1968, pp. 65-70.

5. L.A. Shepp and Y. Vardi, "Maximum Likelihood Reconstruction for Emission Tomography", *IEEE Trans. Nucl. Sci.*, **NS-32**, pp. 3864-3872 (1982).

6. The Infrared and Electro-Optical Systems Handbook, Vol. 2: Atmospheric Propagation of Radiation, Fred G. Smith, ed., ERIM and SPIE 1993, pp. 114 –120.

INVESTIGATION OF INFRARED SEEKER PERFORMANCE AGAINST GROUND TARGETS IN CLUTTER

for the Proceedings of the
Ground Target Modeling and Validation Conference

MeLisa B. Cannon, Ronald C. Passmore
Missile Research, Development, and Engineering Center
U.S. Army Aviation and Missile Command

Mark Chambliss, Carrie Kimbel, Daniel Konkle
Dynetics, Inc.
Huntsville, AL

ABSTRACT

The Aviation and Missile Command (AMCOM) Missile Research Development and Engineering Center (MRDEC) has been investigating the performance of tactical infrared seekers against ground targets embedded in clutter. Figures of merit in this investigation include detection performance utilizing a generic auto-detection routine, a gray-level co-occurrence matrix track metric (GLCM-TM), 'Model A' clutter characterization, and tracker performance from three types of trackers. State-of-the-art infrared imagery from a captive flight test was utilized which consists of dynamic imagery in five-second segments at various ranges and environments including both stationary and moving targets. Previous evaluation of this data included man-in-the-loop (MITL) performance using fifty civilian observers, which will be compared with these non-MITL performance results. The detection routine and trackers used in the evaluation, though believed to be representative of current threat capabilities, are not program specific but demonstrate only generic capabilities. The purpose of this study was to evaluate trends between a clutter metric (Model A), auto-tracking prediction (GLCM-TM), and actual detection and tracker results. In addition, these results will be compared with observer performance. This paper will describe the auto-detection routine, the generic trackers chosen for evaluation, and the clutter and tracker figures of merit. Trends and comparisons will also be shown to illustrate how these characterizations can influence infrared seeker performance model improvements. All of these investigations are part of MRDEC's effort to understand and properly model low signature targets embedded in clutter.

INTRODUCTION

There are many ways to investigate seeker performance on an analytical level, but generally can be divided into two categories: MITL and non-MITL. Both of these areas are valid for seeker performance analysis depending on the system being studied. In FY98, an observer experiment was conducted utilizing unclassified imagery of ground targets embedded in clutter. Specific runs were identified for their difficulty; only those with visibly low signatures were analyzed. Over fifity observers were used to determine detection within twenty six runs of five second dynamic scenes taken from a captive flight test (CFT). Observer and modeled results were published for MITL analysis[1]. Another important analysis excludes the human observer; therefore, non-MITL performance and metrics were used in MRDEC's latest analysis which is the subject of this paper. This analysis falls in two general areas: background analysis (Model A and GLCM-TM) and non-MITL performance analysis (auto-detection and tracking). The following sections will describe the methodology and results of each of these analyses.

BACKGROUND ANALYSIS

Clutter background analysis is a particularly important portion of performance evaluation for IR seekers. Detection performance is clearly a function of clutter level as it effects not only the signal to clutter ratios but also competes with the texture within the target itself. Two analyses were performed with the Gauss-Markov Model A metric and with the Gray Level Co-occurance Matrix Track Metric. These two metrics differ in one important aspect: target texture consideration. Model A is a two dimensional power spectral density representation of clutter –

only clutter. Whereas the GLCM-TM is a measure of how closely the target structure matches the background structure. Model A has been used with reasonable success for air defense seeker applications. With this application, the target structure is of little consideration as detection normally occurs with point sources or marginally extended sources (few pixels on target). The following two sections describe briefly the two types of background metrics and the results from the observer experiment.

Gauss-Markov 'Model A' Metric

The Gauss-Markov Model A clutter model characterizes clutter with a two dimensional power spectral density fit to represent various background types. Figure 1 depicts the calculation of the clutter PSD with this method using three parameters including global standard deviation in background temperature, clutter correlation length, and power roll-off. A more detailed description of Model A and its origin can be found in Reference 2.
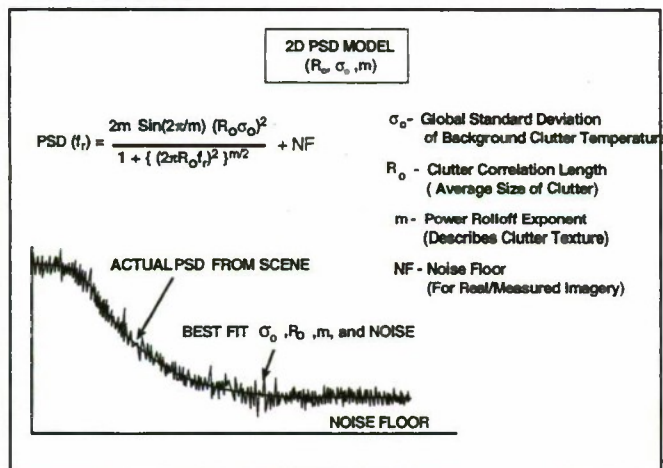
Figure 1. Gauss-Markov Model A

The PSD of the clutter is integrated with the modulation transfer functions (MTF) of the system and signal processing to determine the residual clutter term denoted as σ_C and shown in equation 1. These parameters can be extracted from specific images or refenced from a clutter library by background type.

$$\sigma_c{}^2 = \iint MTF^2{}_{SYS} * MTF^2{}_{SP} * PSD_{CL}\, df_x\, df_y \quad (1)$$

Residual clutter, σ_C, was calculated for each of the twenty six runs and were compared with the

observer performance to determine trends. Figure 2 shows this comparison.
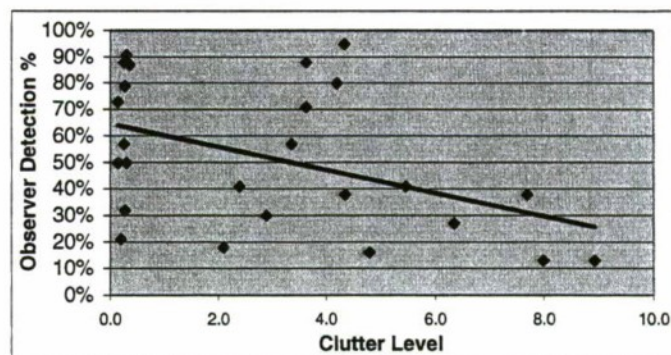
Figure 2. Clutter vs Observer Detection

Though the clutter appears scattered based upon observer detection, the general trend (shown with the line in Figure 2) appears to be as expected, the higher residual clutter levels result in lower detection. Reduction of scatter is also noticed for higher clutter levels suggestiong that in lower clutter levels, observers may leverage from other variables to make detections until the clutter level gets so large the target becomes difficult to distinguish and visual cues are suppressed.

Grey Level Co-occurrence Matrix (Track Metric)

The other background calculation investigated in this study was the Grey Level Co-occurrence Track Metric (GLCM-TM). The Track Metric measures the background and target structure and intensity levels in each image. It is strongly dependent on image, target, and aspect. Therefore calculations were made for each image sample. A detailed description of the GLCM-TM methodology has been previously reported [2,3] and will not be repeated here. However, the general concept is that the track metric compares spatial structure and intensity for both background and target. If target structure and intensity levels are very similar, the TM will reflect a low number, or low trackability measure. Because the GLCM-TM was originally intended as a replacement to delta temperature as a measure of required signature for acquisition, a trend was expected with man-in-the-loop detection because of its image-based dependency as humans also require.

Figure 3. Track Metric vs. Observer Detection

However, Figure 3 shows that very little trend exists between the two for this experiment data. This could perhaps be attributed to the way the data is processed for the track metric vs the observer. In the experiment the observer had to search the FOV, but given the training prior to the test, had an idea of the "general" location of the target. The track metric, however, was calculated knowing the x,y coordinates and took into account only the most local background (2x the target gate size). Perhaps more consistent results would have been made with the observers if a larger background was considered in the calculation.

One would also expect a level of dependency with clutter level. Figure 4 does show a trend as expected: higher clutter levels result in decreasing trackability (low track metric level).
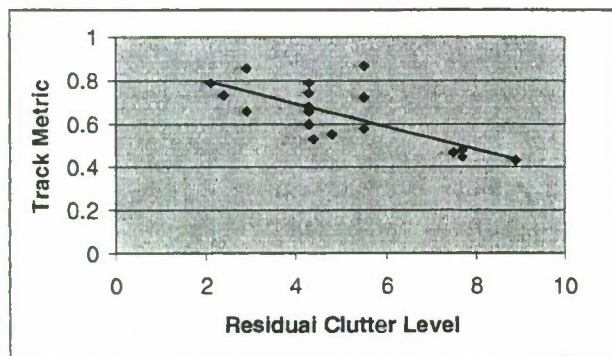
Figure 4. Track Metric vs Model A Clutter

Only the stationary target runs are included in Figure 3. The track metric is not dependent on moving targets whereas the residual clutter levels using the method in equation 1 are considered through the MTF calculation for MTI processing. When MTI processing is performed, only a minimal amount of residual clutter can be seen which is essentially constant for the moving target runs whereas the track metric varies the same as stationary type runs.

PERFORMANCE ANALYSIS

In this study, seeker performance was calculated in two ways: auto-detection and tracking. The auto-detection, not to be confused with automatic target recognition (ATR), is a fairly simple routine using common spatial filters to enhance target-like features in the image, and utilizing a threshold to produce detections. Three baseline tracker algorithms were also utilized to determine if detected, the trackability of the targets in clutter backgrounds. The trackers chosen for this analysis included the hot spot (HS), Bayesian (BS), and moving target indicator (MTI). The following sections describes each routine and the results based upon the observer experiment data.
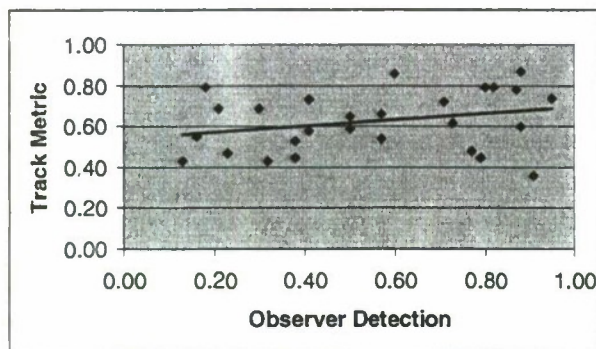
Auto Detection

The auto detection routine is used in two parts: clutter threshold (clutr) and auto-detection (detr). The clutr routine, as the name suggests, samples surrounding clutter to determine a table of gray-level thresholds versus number of detections (true and false). Therefore, the user is allowed to choose a threshold for the image(s) based upon the desired false alarm rate using a series of images that calcualtes the range of thresholds vs number of detections. Using the selected threshold, the detection routine (detr) is used on every frame of imagery. Detections are shown on the imagery by colored dots, with the primary detection in red. The user is then prompted to enter if it is a true or false detection via three possible answers: target, clutter, or other. The results from the user are logged in a file.

For this study, thresholds were selected at approximately 4% false alarms, consistent with current tracker capabilities. A center region of 100x100 pixels was chosen to analyze as the region of interest (ROI). Only in cases in which the target was not within the center 100x100 region that an alternate 100x100 region were selected to ensure the target remained throughout the sequences. It was impractical to process the entire field of view due to the diverse backgrounds and lack of sophisticated routine. Using the selected threshold, the detection routine (detr) was processed over the ROI. Each pixel that exceeded the chosen threshold was shown by color on the image. The user determines if any of the exceedences is a true detection. Only the first

detection of each sequence was noted, assuming that once a detection was made, a tracker would be implemented (not continued detection routine). A detection percentage was calculated by the first correct detection frame over the total number of frames. For example if the target was detected on the first frame, that sequence was given a score of 100% detection. If the first detection was on the 75th frame (out of 150 possible) the score for that sequence was 50%. This seemed to be a fair assumption given that the false alarm rate was kept constant for every sequence.

Given the above assumptions, the results to the auto-detection routine are shown if Figures 5 and 6. The first of these figures shows the comparison between automatic detections and clutter level. Figure 4 shows the trend between clutter, track metric, and auto-detection.
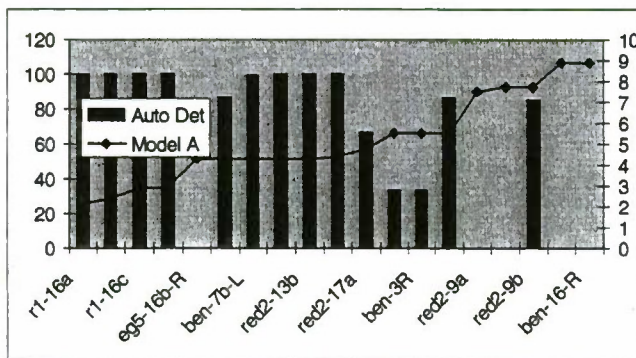
Figure 5. Auto Detection vs Clutter Level

A trend can be seen that consistent detections are made with the lower clutter numbers, whereas the higher clutter numbers (on right side of chart) result in sparse detections. The moving target sequences were not included since the clutter level was essentially constant due to the MTI processing.
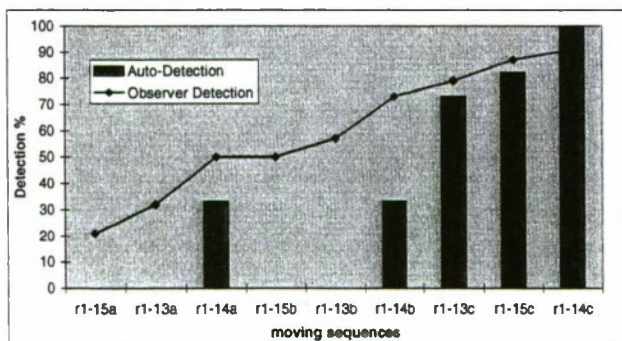
Figure 6. Moving Target Analysis

Though no specific trend was observed between the auto-detection and observers for stationary

targets, the moving target analysis was as expected. There was gradual degradation in performance for the observers who naturally use movement as cues. The auto-detection which did not consider movement only had consistent detections on the closer range sequences (denoted with 'c' – right side of Fig 5). It was also observed that the higher signal to noise ratio requirement for the auto-detection routine was not as robust as the multiple cues used by the observers. A similar gradual degradation for observer performance was seen for the stationary runs with sporadic high detections with the auto-detection. The lack of detections by the detr routine was mostly caused by warmer clutter in the scene that essentially "threshed out" the target. A solution to this problem might be improving signal processing before executing the detection routine.

Tracker Performance

The proposed tracking algorithms consist of three levels of tracking to represent low, moderate, and high levels of algorithm sophistication. The three algorithms will be referred to as the hot spot (HS), Bayesian (Bay), and moving target indication (MTI) trackers. These algorithms are listed in order of their sophistication level from low to high. The sophistication levels are relative to each other and are not intended to classify these algorithms with respect to any algorithms outside this effort. These tracking algorithms are available in the Infrared Seeker Algorithm Tool (ISAT) which was used to determine the performance of each algorithm against the image sequences [5]. The tracking results are presented as number of tracked images throughout the sequence. There was no attempt to assess automatic target acquisition algorithms; therefore, it was assumed that the target location, size, and shape were provided to the tracking algorithms. Also, fielded systems would likely contain algorithms with additional features and more specialization than those available in ISAT. For example, ISAT does not provide track gate (TG) optimization algorithms, so the TGs were optimized for each sequence by user intervention. Other tracking parameters were optimized by the user to maximize performance. An overview of tracker implementation and results follow. Detailed descriptions of all the algorithms may be found in Reference 4.

186

Hot Spot Tracker

The track gate of the HS tracker was centered on the target and its dimensions were held constant throughout the image sequence. The maximum per-frame motion encountered in the sequence determined the size of the track gate. The movements of both the captive-carry platform and the target caused this motion. If the target and image motion combination was less than the smaller dimension of the target, the track gate could be set to the size of the target or, in some cases, even smaller. If the target and image motion combination was larger than the target in either direction, the track gate had to be larger than the target to prevent the target from moving outside of the track gate in the following frame. The HS tracker updated the location of the center of the track gate to the location of the highest intensity pixel inside the initial track gate location. In all the evaluations, the track gate size and location were optimized for the best tracker performance via user inputs.

Baysian Tracker

The Bay tracker classifies pixels based on their similarity to statistical models for the target and background. The TG of the Bay tracker has two components, the target pixel gate (TPG) and the background pixel gate (BPG). To initialize the tracker, the TPG was centered on the target and matched to the size of the target. The BPG was also centered on the target but was larger than the size of the TPG. The actual size of the BPG was based on the optimal performance for a given image sequence. The BPG excluded the area designated by the TPG. The target statistical model was formed based on the pixels inside the TPG. The background model used only pixels inside the BPG. The mean, standard deviation, and correlation coefficient were calculated for the TPG and BPG areas for the first frame to determine the bivariate normal distributions for the target and background pixels. On the second frame, each pixel within the TG was classified according to its probability of being either a background pixel or a target pixel based on Bayes' law. The TG was then centered on the new target pixels and the statistics were updated. Calculation of the correlation coefficient required the user to specify the offsets between the current pixel and its statistical pair. In all cases the size and location of the TPG and BPG were optimized as well as the statistical pair offsets.

MTI

MTI algorithms attempt to enhance the target-background contrast by first aligning, and then subtracting, two frames of an image sequence. Since the two images are aligned with respect to stationary background clutter features, the clutter is greatly reduced by the subtraction. Objects that are moving relative to the background will not coincide after the image alignment and tend to be reduced less. Theoretically, the only pixels remaining after the subtraction are those which fall on moving targets. The MTI tracker was not as sensitive to the TG size as the previous two tracker methods. The TG of the MTI tracker was centered on the target and sized to be much larger than the target. The location of the TG was updated on each frame by centering it on the maximum intensity pixel within the TG. Because the targets in the observer data moved slowly, frames were skipped to increase the apparent per-frame target velocity. For example, frames 1 and 3 may have been registered pairs as opposed to frames 1 and 2. Using this option in the current implementation in ISAT, frames 1 and 3 are paired then frames 3 and 5 are paired and so on. From a tracking standpoint, this caused the data rate to be lower than the frame rate. Regardless, successful tracks at lower data rates should also be successful at higher rates.

Tracker Results

Twenty-two out of the twenty-four analyzed targets were tracked throughout the entire sequence by at least one tracker (Figure 7). Only one sequence was not tracked by any method. The MTI tracker was successful in every sequence that contained a moving target. The HS tracker was successful in 22 out of the 24 image sequences. The track gate of the HS tracker was sized as small as possible based on the amount of per-frame image motion in the sequence. Because the track gate of the HS tracker was small, it helped prevent opportunities for the tracker to be confused by clutter. The Bay tracker performance was worse than the HS tracker. However, its performance was degraded due to excessive image motion and the lack of gimbal or IMU data. Also, dynamic gate management algorithms would increase the performance of the Bay tracker. It is not surprising that if the target location is correctly designated and the gate sizes are optimized for each sequence, even nominal trackers such as thes used in this study perform very well.

187

Seq	Hot Spot	Bay Trkr	MTI
red1-13a			
red1-13b			
red1-13c			
red2-5a			N/A
red2-5b			N/A
ben-3-L			N/A
ben-3-R			N/A
red1-15a			
red1-15b	30%	30%	
red1-15c			
red2-9b		40%	N/A
ben-7b-R			N/A
ben-7b-L			N/A
red1-16a			N/A
red1-16b			N/A
red1-16c			N/A
red1-14a			
ben-16-L		57%	N/A
ben-16-R			N/A
red2-17a			N/A
red1-14b			
red2-13b		50%	N/A
red2-17b			N/A
red1-14c			

Legend:
- Successfully Tracked Entire Sequence
- Trk% Successfully Tracked Partial Sequence
- Did not track target
- Not Applicable (stationary target)

Figure 7. Tracker Results

The track metric was also compared with the tracker results (Figure 8). There appeared to be no correlation between the two. This result is thought to be attributed to the optimization that occurred for the tracker for each sequence. Future efforts will include running the trackers for every sequence without optimization using some default mode. Comparisons will be made then as to if the sequences with lower track numbers required the most optimization to achieve optimal results.
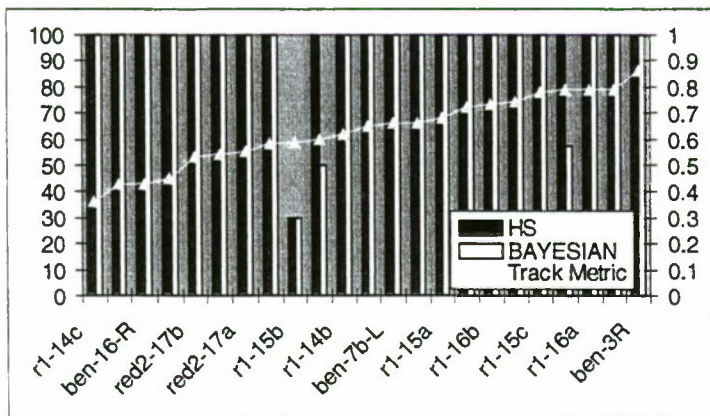
Figure 8. Track Metric vs Tracker Performance

SUMMARY

This paper has presented the most recent activities in a continuing effort to improve the capability of IR seeker performance models. This thrust is directed toward more accurately modeling the IR imaging seeker capabilities for resolved, low intensity targets embedded in complex clutter. Both MITL and non-MITL capabilitiees are being considered and the progress of this activitiy is coordinated with other efforts invovled in reducing this complex modeling issue into a few parameters.

Several very diverse types of analysis were conducted on the observer experiment imagery. Background analysis was made observing trends with the Model A clutter and the GLCM track metric. The residual clutter parameters derived from Model A did show a trend compared with observer performance as higher clutter produced decreased detection probability as well as showed less scatter than with lower clutter levels. A noticable trend was not expected dut to the metric's lack of target structure consideration. Future studies will investigate Model A's applicability in performance model for extended sources (e.g., ground targets). Though a trend could be seen with the track metric and clutter level, it was disappointing that no trend could be shown with the observer detction and track metric. Modifications are planned and future comparisons will be made with non-optimized routines and possibly more observer detection data.

The auto-detection routine did not demonstrate a state-of-the-art sophisticated algorithm, but showed the importance of SNR on detection capability when MIL is not used. In general, MIL detection was better than auto-detection as expected since the routine does not emmulate the many signal processing roles of the human brain. Tracker algorithms, however, prove very robust if provided target information. The real issue appears to be in search and accurate detection also in plans for near term investigation.

The results presented here are but a step toward the complete understanding of this complex modeling challenge. Similar efforts will continue under a variety of AMCOM sponsorhips to investigate and determine meaningful background clutter metrics and formulations of target signatures that are conducive to determine IR

seeker performance. Additional analytical work is planned and experimental efforts will also provide additional data from structured ground and captive flight tests.

REFERENCES

[1] Cannon, R., Jones, S., Kusner, R., "Observer – Based Experiment for Man-in-the-Loop Target Detection Using CFT Seeker Imagery", Proceedings from Ground Target Modeling & Validation Conference, Aug 1998.

[2] Dynetics Inc., "IR Clutter Model Handbook", October 1996.

[3] Brackney, B. A., "Improved Trackability Metric for Resolved Targets," Dynetics, Inc., Memo-SAE-95-006, Mar 1995.

[4] Chambliss, M., "User Guide for Infrared Seeker Algorithm Testbed Version 2.0," Dynetics,Inc., Memo-SAE-98-001.

[5] Chambliss, M., and Olree, K., "Performance of ISAT Algorithms Against Observer Experiment Image Sequences", Technical Memo: SAE-99-001, 15 April 1999.

Airborne FLIR Performance Prediction and Optimization

S. S. Yoo, H. Tran

Northrop Grumman Corp.
600 Hicks Rd
Rolling Meadows, IL 60008

ABSTRACT

In recent years, a number of critical advancements have been made in infrared detector material preparation, detector fabrication technology, and readout integrated circuit (ROIC), as well as cryogenic technology. These combined technologies allow production of higher resolution IR focal plane array (FPA) of superior quality, such as 512x512, 640x480, or larger pixel format FPA devices of either HgCdTe, InSb, or QWIP. In general, high resolution FPA devices are needed for infrared imaging cameras that provide long-range detection of small objects, particularly in military applications.

In this study, the performance of an airborne infrared camera with large pixel format FPA forward looking IR (FLIR) is predicted for various optical configurations. For airborne imaging, the space constraint for optics and electronics becomes crucial, as does the limited optical aperture size. The first constraint will affect the focal length of the optics, hence the instantaneous field-of-view (IFOV); the latter will affect the F-number of the optics, hence the sensitivity of the system. Both constraints will critically affect overall infrared camera performance. Compromise is essential to achieve the best performance of an airborne FLIR with a given line-of-sight instability. Through this study, we exercised the optimization process for long-range ground target detection and conducted a trade study between Mid-wave IR (MWIR) and Long-wave IR (LWIR) FLIRs.

INTRODUCTION

The digital modeling tool FLIR92 by NVESD has been the standard means to predict the performance of FLIR cameras; however, this standard has some shortfalls when applied to staring arrays, because its predictions are generally pessimistic for staring array sensors. Very often pilots have noticed that they can recognize targets beyond the Nyquist based predictions. Users have typically overcome the shortfalls of FLIR92 by extrapolation and other such "unofficial" techniques to make the basic model work, i.e., to be representative of actual FLIR range performance[1].

NVTherm[2] is a recently released FLIR modeling tool by NVESD that replaces FLIR92. This new model is more directly applicable to typically undersampled staring arrays and considers performance of FLIRs past their Nyquist frequency. In addition, the combination of higher sensitivities of staring arrays and the contrast limitations of the human eye has recently become a critical limitation of staring array imaging systems that was incorporated in NVTherm.

A beta version of NVTherm was used for system design trade study and performance prediction of large format FLIR in support of on-going next generation ground target pod development at Northrop Grumman, Defensive Systems Division at Rolling Meadows, Illinois. The MODTRAN® AFRL/GL (ONTAR Corp.) atmospheric transmission model and ACQUIRE algorithm were also used in support of performance prediction. (*ACQUIRE Range Performance Model for Target Acquisition Systems* is a NVESD software program, normally used with FLIR92). An IR sensor radiometric model was developed for various types of IR imaging systems. This model, verified with numerous sensor system characterization test results, was also used to increase understanding of advanced imaging systems and, in part, to support the verification of NVTherm performance prediction.

SYSTEM F/#

Many parameters affect the range performance of an IR imaging system. Each parameter has a significant impact

on the FLIR performance. Depending on its application, the system needs to be optimized for its specific applications, such as hot or cold weather background at long or close range. In general, low F/# systems provide excellent sensitivity and are useful as accurate radiometers for mapping temperature differences of an object or geographic surface. Though the sensitivity of the incorporated detector device remains constant, the system F/# affects the system sensitivity and resolution. For low F/# systems, instantaneous field of view (IFOV) is relatively large and pixels are exposed to the large area background corresponding to the IFOV. That increases the scene radiance reaching the FPA pixels, and provides a relatively large signal-to-noise ratio (SNR) for small scene temperature changes. In spite of the excellent system sensitivity, these systems have shortfalls for long-range applications, mainly because of their resolution limitation. In the other extreme, high F/# systems provide excellent resolution, a great advantage for long-range imaging. However, such high F/# systems often drive the system resolution into the optics diffraction limited resolution and require long integration time to achieve SNR.

Assuming 100 % cold shield efficiency, a series of Minimum Resolvable Temperature (MRT) values at 300 K background temperature were generated with variable optics F/# using the NVTherm model. In the model, the integration time was set to maintain 50% storage capacitor well fill in order to compare systems of different F/#. The results are compared in Fig. 1. It is noticed that as F/# increases, MRT curves shift to higher frequency, thus making the system more adequate for resolving small targets at long distance.

In spite of great spatial resolution, there are several factors to limit the performance of such high optics F/# systems: diffraction limited resolution, sensitivity due to the reduced optical power for a pixel, and susceptibility to line-of-sight jitter. The optical aperture size often causes optics diffraction limited resolution, especially where high resolution is required. Usually aperture size in an airborne FLIR is limited by space constraint, due to consideration of air drag at supersonic speed. Therefore, with a given maximum aperture size, system optimization is needed to achieve the required performance.

Another consideration is the increase in integration time to obtain reasonable SNR. 50% well fill integration time increases linearly with system F/#. This is because of reduced FPA FOV, or equally reduced pixel IFOV: a reduced amount of optical power is captured by a pixel. The maximum integration time of a real time imaging system cannot be longer than the frame time (i.e. 33 msec

for 30 Hz or 16 msec for 60 Hz frame rate) and this sets the maximum integration time. NVTherm incorporates the effect of integration time. Fig. 2 shows MRT improves with increasing integration time (or increasing scan efficiency).
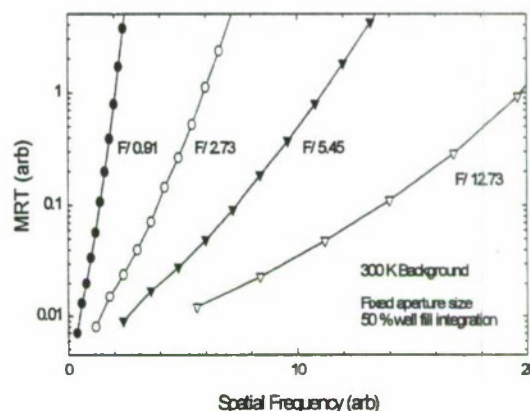
Fig. 1. MRT values generated with NVTherm for various optics F/#.
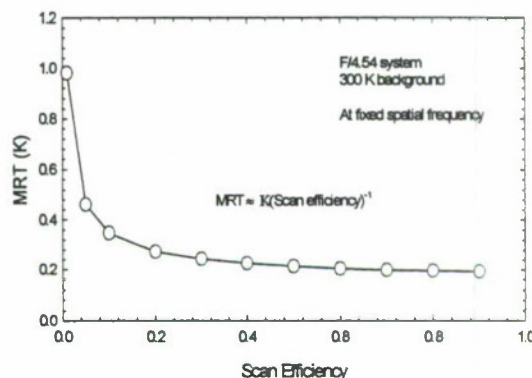
Fig. 2. Scan efficiency dependence on MRT. Note that MRT improves with integration time.

Beside NVTherm, where MRT and modulation transfer function (MTF) are the main objectives, the IR camera radiometric model provides an excellent tool for laboratory system characterization. The radiometric model was developed to support lab testing and predict lab test results. Due to the lack of subsystem component parameters from vendors or manufacturers, very often the specifications of each component are not useful for or are irrelevant to modeling of system performance. Therefore

instead of modeling with subsystem parameters, the final system considered as an unit-under-test is tested radiometrically and the result is compared with model prediction. If there are any discrepancies between the model prediction and the test result, the assumed optics, electronics, and detector parameters can be extracted by fitting the model to the test result. These parameters are then used to compare with manufacturer's specifications. This model has proven to be successful in predicting various test results and is being used to support a number of in-house EO/IR programs. Details of this model will be presented in a forthcoming article.

Based on the radiometric model which has been validated with test results on a number of F/# systems, temporal noise was calculated for the extended F/# range under the assumption of the same sensitivity FPA with a 100% efficient cold shield matched to the optics F/#. In Fig. 3, the calculated temporal noise components were categorized into a photon noise component and a rest-of-noise component called system noise. The system noise consists of FPA multiplexer, electronics, detector dark current, and thermal noise due to optics emission. Notice that the total noise remains the same over the F/# range for 50% well fill integration time. As F/# increases, photon noise decreases and system noise increases. This observation is explained by the fact that pixel irradiance is reduced with F/#, thus resulting in reduced background photon noise. Long integration time mainly leads to increased dark current noise, which is time dependent. Therefore, the longer the integration time, the greater the dark current noise contribution.
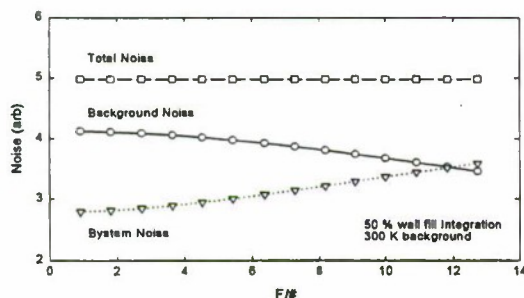
Fig. 3. IR camera model prediction of background noise limited operation. Background photon noise gradually decreases with increase of F/# and becomes comparable with system noise at higher F/#. The total noise of the system remains the same over system F/#.

System responsivity at 300 K background temperature was calculated with system F/# as a variable, and the results are shown in Fig. 4. Irradiance at the aperture is reduced by the solid angle of a pixel, or IFOV, hence decreasing the responsivity at the background scene temperature. It was shown previously that as integration time is adjusted for 50% storage capacitor well fill, total noise should remain the same. As a result, noise equivalent temperature difference (NETD) is found to increase with F/#. NETD is a temporal noise component, which critically affects system MRT of high F/# systems in all the spatial frequencies.
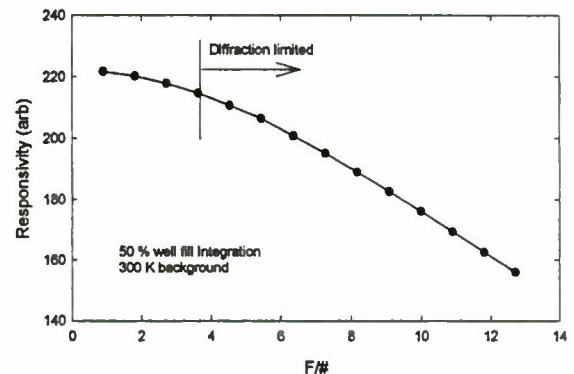
Fig. 4. Responsivity versus System F/#. Because of the reduced pixel irradiance of high F/# system, system responsivity decreases with F/# and, as a result, system NEDT increases.

ATMOSPHERIC TRANSMISSION

MODTRAN was used to calculate the average atmospheric transmission. The selected atmosphere was the mid latitude summer using the standard model options with the rural aerosol model at 5 km visibility. The aerosols, which affect meteorological visibility, are contained in the boundary layer of 0 to 2km altitude in MODTRAN model. The atmosphere is murkiest at ground level and gets somewhat clearer with increasing altitude, and when mixing with higher level air up to the inversion layer, which is at the top of the boundary layer. At an altitude with a given slant angle, the transmission varies with altitude. For example, at 30 degree slant angle, the line of sight travels only 4 km within the boundary layer, mostly affecting the total transmission along the path to the FLIR position. For this reason, the transmission model cannot be simplified with a single transmission coefficient, that could cause underestimated or overestimated transmission assumption.

Still, visibility is not a strong factor for air-to-ground viewing at these angles. It should be noted that on semilog scales the transmission is not an exact Beer's law (straight line) and its use creates significant errors at the longer ranges. It was found that a very good fit (within 1%) over a large range is obtained using a 6-parameter exponential based on the Weibull distribution. Total transmission from a target was computed by MODTRAN along the slant range from 0.5 km with 0.5 km steps at a given slant angle, then the parameters were obtained by fitting the data. Fig. 5 shows the results. This procedure was repeated for all the possible slant angles from 5 to 90 degree for the calculation of detection range.

The impact of LOS jitter to FLIR performance is illustrated in Fig. 6. In this figure, for the given F/4.54 optics with the same IFOV, the degraded FLIR MRT with varying jitter was calculated, then used to extract the recognition range with fixed target size and atmospheric conditions. The recognition range decreases gradually beyond the point where jitter is larger than the IFOV of the FLIR. Well-controlled jitter of less than 50% of IFOV hardly seems to affect the system performance.
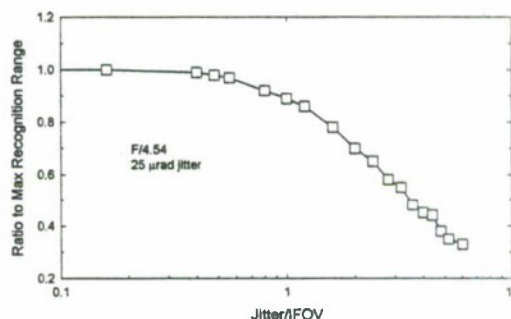
Fig. 6. Degraded FLIR performance due to LOS jitter. Note that when jitter/IFOV is unity, the recognition range is reduced by 15%. Beyond that, the range decreases rapidly.

Fig 5. Comparison of MODTRAN calculation of atmospheric transmission along slant range and a simple Beer's model.

LINE-OF-SIGHT JITTER

The line-of-sight (LOS) jitter due to airbone platform vibration is the most critical factor affecting airborn imaging system design. The effect of jitter is image blur during the integration, which consequently impacts the recognition of small targets viewed on a display. Typically, jitter bandwidth is higher than integration time for high F/# systems, and details at high spatial frequncy become blurred. As illustrated previously for longer range recognition capability, the system magnification should be increased as much as possible. That results in smaller FOV or a higher F/# optics system. Unlike ground systems, LOS stability is greatly affected by the ariborne platform movment and is usually limited by the capability of the stabilization system. For high F/# optics or small FOV, the effect of jitter becomes more critical and results in reduced recognition range.

Fig. 7. FLIR recognition range is affected critically by LOS jitter, especially for long range, high F/# FLIR systems.

High F/# optics systems result in small IFOV, increasing suceptability to line-of-sight jitter and impacting system performance critically. With NVTherm-generated MRT for various F/# systems, the effect of jitter was introduced to calculate a target recognition range under the same conditions. The result is shown in Fig. 7, where the jitter impacts high F/# systems more critically than low F/# systems. Athough high F/# systems seem promising for

long-range target recognition, the presence of jitter severely degrades the system performance.

TARGET MODEL

Typically, armored ground vehicles and bunkers are considered as ground targets for airborne FLIRs. When ground targets are viewed from altitude with a slant angle, the projected target size is larger than the size seen at normal viewing angle.

In order to estimate the target size projected onto the line-of-sight, the target was assumed to be a block having the length (l), width (w) and height (h) dimensions of the target. Then the specified target size was modified by the cosine effects of a non-normal viewing angle and converted to an effective target area. The target is seen as a side view or front view with some increase in height due to the viewing aspect angle.

The following equations were used to compute the projected height of target side and front:

Projected height (side):
$$Hs(\theta) = h \cdot \cos(\theta) + w \cdot \sin(\theta)$$

Projected height (front):
$$Hf(\theta) = h \cdot \cos(\theta) + l \cdot \sin(\theta)$$

where θ is the slant angle of the line-of-sight from the FLIR to the target.

Using the calculated projected heights of target side and front, the effective target size of a typical ground vehicle is illustrated in Fig. 8. It is noticed that the effective target size initially increases with slant angle, reaching its peak at 50 degrees for the target side view (>20% net increase) and at 70 degrees for the target front view (>30% net increase). This observation illustrates that at higher altitude, larger effective target size leads to longer recognition range, or higher probability of recognition at the same slant range.

RANGE PREDICTION

Many environmental parameters that affect the range and performance prediction of an airborne IR imaging system. Each parameter has a significant impact on the final range performance. In addition to the environmental parameters, the target parameters have a significant impact on the range performance of the given IR system[3].

The path by which the range performance prediction was determined is briefly reiterated again. As seen earlier, both atmospheric transmission and apparent target size are dependent on the slant angle of the LOS from the FLIR. For this reason, an approximated transmission by Beer's law may result in underestimated or overestimated range and should not be used unless the range is roughly known. Similarly, for the apparent target size, the variation of the target size is nearly 20% between the size seen at normal viewing angle and the maximum size at a slant angle. This translates to an error of 20% in the range prediction. Therefore, recognition range is dependent on slant angle and the FLIR performance prediction should be carried out in terms of slant angle.
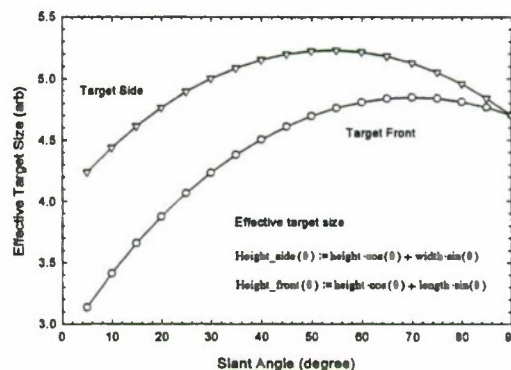
Fig. 8. Effective target size of a typical ground vehicle at slant angles from 5 degrees to 90 degrees.

The predicted range can be calculated by solving the apparent target temperature and MRT curves. The target apparent temperature variation along slant range is obtained by multiplying a given target temperature difference by the transmission. NVTherm-generated MRT is modified with a specified LOS jitter. The spatial frequency of the MRT curve is then converted into range with effective target size and Johnson detection criteria, typically 1 cycle for detection, 3 cycles for recognition and 6 cycles for identification are assumed for 50% probability.

Both the calculated transmission data at a given slant angle and the resulting MRT data are plotted in Fig. 9 to illustrate the procedure. The recognition range can be determined easily by solving these two interpolated curves. This procedure is repeated at the various slant angles to construct the boundary of the required range as

194

shown in Fig. 10. It is noticed that the predicted slant range varies with altitude. Relatively short recognition range is obtained at small slant angles, and the range increases slightly with the angle. This is because of the combination of improved atmospheric transmission and increased apparent target size at the slant angle.
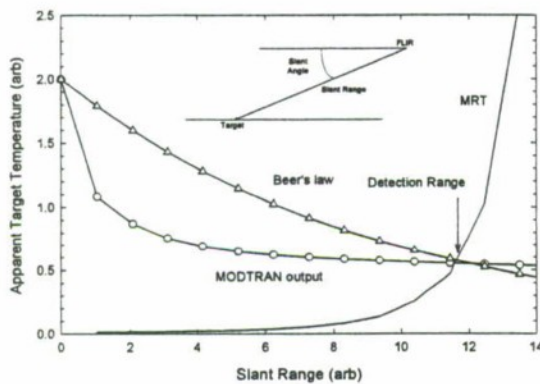
Fig. 9. Recognition range prediction by solving the apparent target temperature and MRT curves.
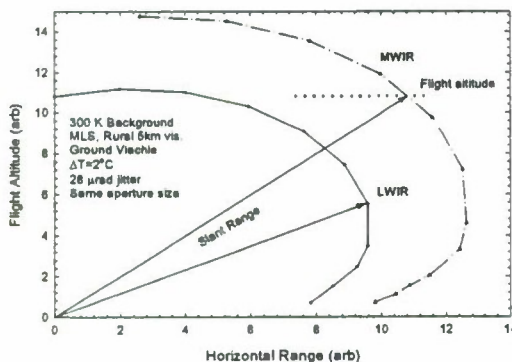
Fig. 10. Predicted recognition range of MWIR and LWIR of the same aperture size. Both FLIR are diffraction limited.

MWIR VS. LWIR FOR AIRBORNE FLIR

Since the introduction of IR imaging systems, there have been many studies and arguments regarding MWIR as a better imaging band than LWIR or vice versa. There are a number of factors involved in FLIR performance prediction and analysis: atmospheric transmission, FPA spectral response, system sensitivity, system resolution,

background temperature, and target characteristics. Depending on specific system application, IR imaging systems operating in either LWIR or MWIR have their own advantages.

For an airborne FLIR, long-range target imaging capability and target recognition/identification is the most important requirement. Most importantly, the system should be capable of providing high resolution imagery. However, the system resolution is often limited by staring FPA detector resolution. In order to overcome the detector resolution, it is a common practice to push the system resolution into the optics diffraction limited resolution. In this case, the aperture size becomes a critical parameter to determine the system resolution. However, for an airborne FLIR there is a physical limitation in the aperture size that is unlikely in ground based systems. In Fig. 11, both detector and optics diffraction limited resolutions are illustrated (4). For this particular MWIR system, diffraction limited resolution occurs at F/2.7: beyond that, the diffraction resolution becomes a limiting factor. It is shown that LWIR optics diffraction resolution is twice as large as that of MWIR for the same aperture size. In order to be comparable with MWIR diffraction resolution, the LWIR FLIR aperture size has to be twice the size of MWIR FLIR.
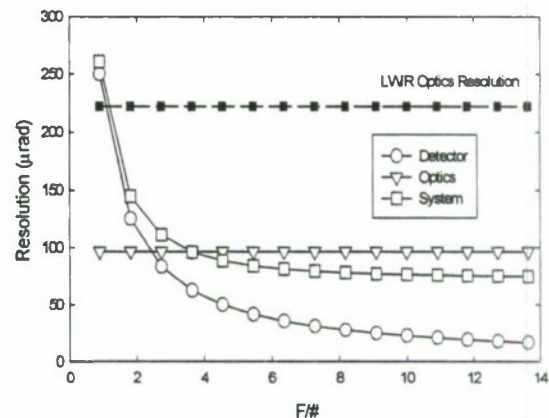
Fig. 11. Comparison of detector resolution and optics diffraction resolution of MWIR to LWIR optics diffraction resolution.

Although more radiance is available from terrestrial background in LWIR than in MWIR, radiance difference due to the temperature difference is higher in MWIR than in LWIR, providing better scene contrast. Furthermore, such large available radiance seems to be often

detrimental to FPA interface to large format ROIC. As FPA format size grows, the integration capacitor tends to shrink because of reduced ROIC cell size. This would be a great concern for LWIR FPA development toward a large format unless the system spectral bandpass is not narrowed.

Atmospheric transmission depends on the amount of water vapor, aerosols, and molecules present in the air. Aerosol affects MWIR transmission; water vapor affects LWIR more than MWIR. Typically, in good visibility and at long meteorological ranges, MWIR transmission is often better. Furthermore, by looking down from high altitude, the transmission is effected in the boundary layer near the ground. Beyond the boundary layer, the MWIR transmission is as good as or sometimes better than that of LWIR.

SUMMARY

Using NVTherm and the Acquire algorithm, airborne FLIR performance was predicted. In order to achieve long-range target imaging, it is desirable to design a system that provides high resolution and magnification. High resolution, obtained by high F/# optics systems, results in diffraction limited resolution. Some degree of diffraction limited resolution is beneficial to improve the system resolution. However with higher system F/#, the system resolution tends to become limited by detector resolution, and the system responsivity is reduced rapidly as well because of the reduced IFOV. Although the system resolution is improved with high F/#, the range performance is critically affected by LOS jitter. It was found that jitter equal to the IFOV results in 10% reduction in recognition range.

Unlike ground-based system performance analyses, the variation of the atmospheric transmission and the apparent target size should be considered. The apparent target size was shown to be slightly larger at a slant angle, and increases with the viewing slant angle or higher altitude, resulting in longer predicted recognition ranges. It was also shown that the atmospheric transmission along a slant angle should not be assumed as a fixed average transmission because of altitude dependent transmission.

Finally, spectral band selection between MWIR and LWIR was discussed in terms of system resolution and atmospheric transmission. Under optics diffraction limited resolution conditions, the system resolution depends on the optics aperture size and IR band wavelength. With a given aperture size, MWIR provides

twice as much resolution than LWIR. It was also shown that the atmospheric transmission of both MWIR and LWIR along a slant range is slightly affected by meteorological visibility within the boundary layer.

REFERENCES

[1] Wittenstein, W., "Thermal range mode TRM3". SPIE Proceedings Vol. 3436 "Infrared Technology and Applications XXIV." 1998.

[2] Vollmerhausen, R., Driggers, R.G., "NVTherm: Next Generation Night Vision Thermal Model", Proceedings for 1999 Meetings of the IRIS Specialty Groups on Passive Sensors, Vol 1, p121 (1999)

[3] Holst, G. C., "Electro-Optical Imaging System Performance". JCD Publishing and SPIE Optical Engineering Press, p296 (1995)

[4] Holst, G. C., Photonic Spectra, p144, Jan (1999)

ACKNOWLEDGEMENT

Thermal Modelling of Target and Background

Eivind Strømman
Norwegian Defence Research Establishment
PO Box 25, N–2027 Kjeller, Norway

ABSTRACT

Mathematical models and an image manipulation system have been developed for use for camouflage evaluation, thus reducing the need for costly field trials. The aim of the work has been to present an image of a camouflaged object in a realistic terrain background for any given weather condition. The work has been pursued along two directions:

1) Modelling of buildings and bare ground based on extensive measurements of all influencing parameters like air temperature and humidity, ground temperature, wind speed, in- and out-going radiation, sun position, etc. Other factors included in the modelling were heat convection and material properties. The modelling comprised even camouflage nets over the buildings. The results from the models – apparent temperatures – were verified by radiometric measurements.

2) Inserting images of real objects into images of real backgrounds, thus make it possible to adjust the object/background contrast by camouflage measures like screening of hot parts, using camouflage nets and even IR smoke. The smokes were represented by smoke grenades modelled in a dynamic way showing the diffusion of the smoke and the influence of the wind.

Work is in progress to model background vegetation of various types, to model some military objects, and to fuse the results into a common image. The modelling will be combined with measurements of all relevant parameters and with recording of IR imagery at two stations over a year in Norway.

INTRODUCTION

In the mid 70–ties we worked on camouflaging a Norwegian NIKE site both in the visual and thermal spectral band. During this work we started modelling the thermal signature of the various buildings at the site. The work was later continued with modelling of characteristic Norwegian background types and relevant camouflage measures. The result of this work was a simple, but very effective model for predicting surface temperatures.

In our work on studying camouflage effectiveness, field trials was extensively used. In order to reduce costs, an effort to use image manipulation as a substitute, or rather a supplement to field trials, was started. This work was done in cooperation with the pattern recognition group at our establishment, and resulted in a system that made it possible to vary contrasts and introduce various camouflage measures in both visual and thermal images.

The work on mathematical modelling and synthetic images are continued as we are lacking good mathematical models for vegetation like trees and bushes and for military vehicles.

THERMAL MODELLING

Concrete wall

The simplest object to model is a concrete wall, and the principles used here is the foundation for all further modelling. The model is described in reference [1] and it handles both the heat exchange with the environment and the heat flow through the wall (Figure 1).

Absorbed direct solar radiation	$\alpha \cdot \varphi \cdot \tau \cdot \cos u \cdot \cos v$
Absorbed scattered radiation	E_s
Absorbed heat radiation from sky and terrain	$\frac{1}{2}(\varepsilon_0 + 1) \cdot \varepsilon_1 \cdot \sigma \cdot T_0^4$
Emitted heat radiation	$-\varepsilon_1 \cdot \sigma \cdot T_1^4$
Heat convection	$h \cdot (T_0 - T_1)$
Heat conduction	
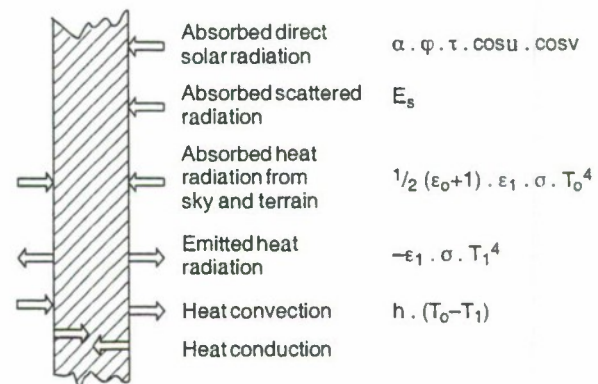
Figure 1 Thermal processes

where the parameters are:

α wall's absorptance of solar radiation
φ solar constant
τ atmospheric transmittance
u, v sun elevation and azimuth angles respectively
$\varepsilon_0, \varepsilon_1$ emissivity of atmosphere and wall respectively
σ Stefan Boltzmann's constant
T_0, T_1 air and wall surface temperatures respectively
h coefficient of convective heat transfer (function of wind speed)

The modelling of the heat conduction through the wall is based on the assumption that the extension of the wall is much greater than its thickness, thus the problem is reduced to calculation of a one dimensional heat flow.

Consider a row of unit area slabs of thickness Δx through the wall with centres 1, 2, 3 etc as shown i figure 2. At the wall surfaces there is only half a slab, so that the centre lies on the surface.
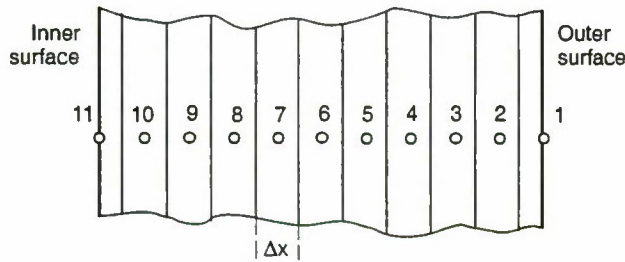
Figure 2 The slab row of the wall

At a given moment, t, the temperature at the various slab centres are denoted T_1, T_2, T_3 etc and after a small time interval Δt: T'_1, T'_2, T'_3 etc. The heat change in one of the inner slabs, for instance slab 2, during the time interval Δt will be

$$\rho \cdot c \cdot \Delta x \cdot \left(T'_2 - T_2 \right) \tag{1}$$

where ρ is the mass density of the wall and c is the wall's specific heat capacity. The heat exchange is a function of the thermal conductivity of the wall, λ, and of the temperatures of the neighbouring slabs. As a first approximation the two neighbouring slabs are considered, and the heat balance of slab 2 is obtained by

$$\rho \cdot c \cdot \Delta x \cdot \left(T'_2 - T_2 \right) =$$
$$\Delta t \cdot \left[(T_1 - T_2) + (T_3 - T_2) \right] \cdot \lambda/\Delta x \tag{2}$$

For the two outer slabs the external heat exchange with the surroundings has to be added (figure 1), and the equation for slab 1 then becomes

$$^1/_2\rho \cdot c \cdot \Delta x \cdot \left(T'_1 - T_1 \right) =$$
$$\Delta t \cdot \left[(T_2 - T_1) \cdot \lambda/\Delta x \right.$$

$$+ a \cdot \varphi \cdot \tau \cdot \cos u \cdot \cos v + E_s$$
$$+ ^1/_2 \left(\epsilon_o + 1 \right) \cdot \epsilon_1 \cdot \sigma \cdot T_o^4 - \epsilon_1 \cdot \sigma \cdot T_1^4$$
$$+ h \cdot (T_o - T_1) \right] \tag{3}$$

For the inner wall surface there will be a similar equation, but without the terms referring to solar radiation (third line of equation 3.

Input parameters in addition to material constants are: geographical coordinates (for calculating sun position), air temperatures and wind speed. As start condition for the calculations is used a very rude estimate of wall surface temperature just before sunrise when the heat balance is in equilibrium.

After several runs it has been found that time intervals of $\Delta t = 30$ sec between each calculation are small enough to give reasonable results for a concrete wall 30 cm thick and divided into 11 slabs as shown i figure 2. Further, it has been shown that increasing the number of slabs will not increase the accuracy of the results significantly.

Composite walls

Most walls are, however, more complicated than the simple concrete wall. The normal structure of a wall is two outer hard, weather resistant surfaces with an insulation material in the middle. These different wall layers have their own specific material constants, thus making the modelling a bit more complicated.

In principle the modelling of a composite wall is the same as for the simple concrete wall. Each layer in the wall is divided int 11 slabs, and the heat transfer for the centres situated on the adjoining surfaces of two neigbouring layers are calculated as before – described in reference [1].

In figure 3 are given examples of two types of walls: The wall of a wooden building (wooden outer surfaces with a layer of insulation in between) and the wall of a metal van (thin metal plates with a hard type of insulation). The modelling of the van wall is simplified by assuming that the thermal conductivity in thin metal plates is high enough to give a practically no heat gradient through the plates.

Calculation of the surface temperature of two such walls has been performed in parallel to measuring the surface
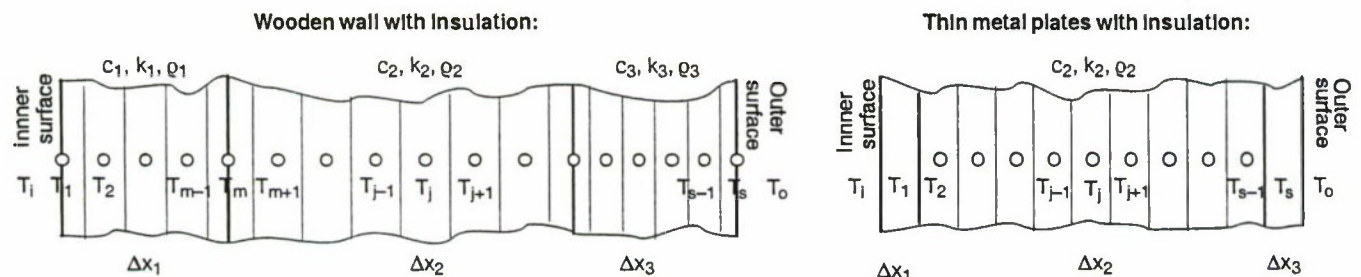

Figure 3 Slab rows of composite walls

198

temperatures. Some results are presented in figure 4. The measuring site was in Southern Norway and the date May 6. The wall orientations is shown in the figure, and the walls were all olive green coloured. Inner temperature of the buildings were 20°C and the wind velocity about 2–3 m/s. The measurements were undertaken with a Barnes PRT–5 radiometer.
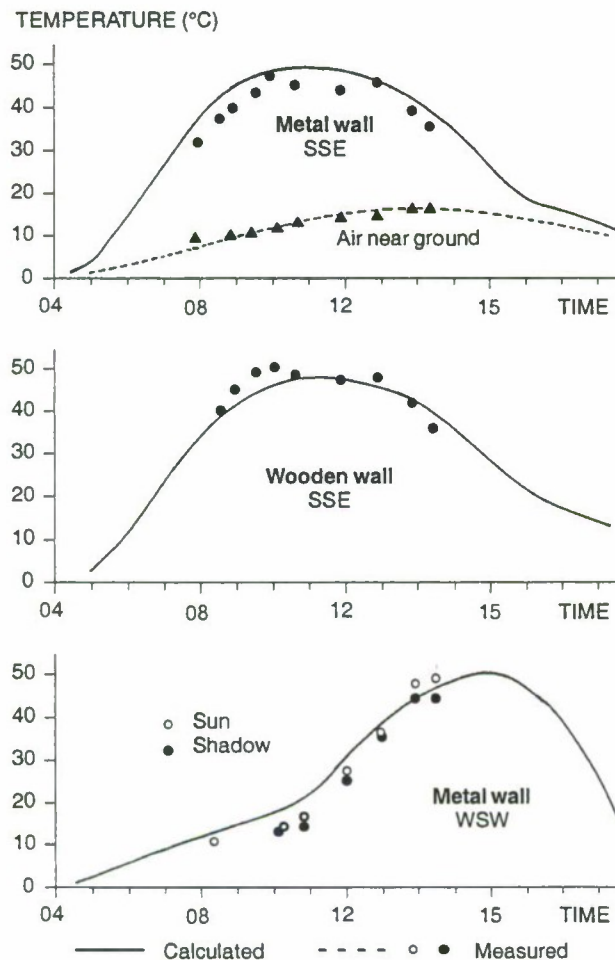
Figure 4 Surface temperatures of composite walls

As can be seen from the figure, most of the differences between measured and calculated temperatures for the composite walls lies well in between ±5°. The results also clearly reflect how the surface temperatures are dependent of wall orientations.

Net camouflaged building

Applying a camouflage nets to a building as shown i figure 5, gives two obvious important effects on the apparent temperature measured with a radiometer:

- When viewed from outside, the measured radiation from the camouflaged object is a mixture of radiation
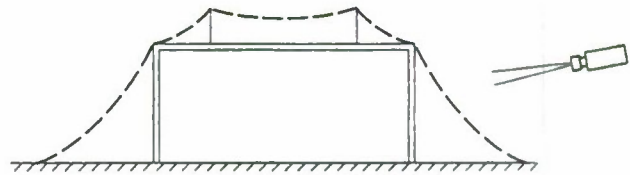
Figure 5 Net camouflaged building and radiometer

from the net itself and radiation from the wall underneath as seen through the openings in the net.

- The sun heating radiation of the wall is reduced to the amount of radiation going through the openings in the camouflage net.

As a radiometer receives radiation from both the net and the wall through the holes in the net, the overall temperature of the camouflaged object can be calculated from the simplified relationship

$$T^4 = f \cdot T_n^4 + (1 - f) \cdot T_w^4 \qquad (4)$$

where f is the fraction covered by the net, and T_w and T_n are the surface temperatures of the wall and net respectively.

When calculating the wall temperature, the radiation exchange with the camouflage net has to be considered. And the heat exchange with the environment has to be modified due to the geometry formed by the net. This means that most lines in equation 3 have to be modified and a few new lines added as described in reference [2].

The camouflage net has a complicated structure which is not straightforward to model. However, the same model as used for walls are found applicable when taking into account the special geometry of the net. Measurements has shown a nearly instantaneous temperature adjustment to sunshine, indicating the heat transfer in the net material can be approximated to a steady-state problem.

The most complicated phenomenon to model is the heat convection (figure 1) which has been handled in detail in reference [2]. The coefficient of convective heat transfer, h, is a function of several parameters of which the typical size of "net leaves", l, and wind velocity, w, are the manageable parameters. Assuming turbulent forced air flow around the net leaves, the coefficient of convective heat transfer becomes

$$h \simeq 6.5 \frac{w^{0.8}}{l^{0.2}} \qquad (5)$$

The camouflage net has an additional effect in that the wind velocity at the wall underneath the net is becoming reduced. Measurements have shown that it is reduced to a fraction 0.4–0.6 when the wind velocity outside the net is 3–5 m/s.

The model has been applied to a construction like that in figure 5, and radiometric measurements have been undertaken on three different days. The results presented in figure 6 show that there is a good correlation between calculated and measured temperatures. On each day the three types of walls presented earlier, were included. The deviations between calculated and measured temperatures are typically less than 5°C.
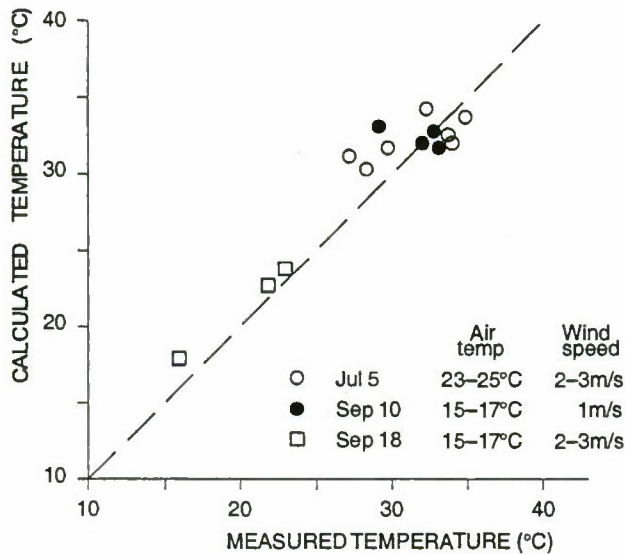
Figure 6 Surface temperatures
of net camouflaged walls

Bare ground

The model can as well be applied to horizontal surfaces as to vertical walls, and it has been modified for bare ground – reference [3]. A sand ground can have layers with different thermal properties due to different water content. And there can be different layers of soil quality. In principle the calculations are quite similar to that of a composite wall. The model has, however, been extended to take into account evaporation of water: the cooling effect and changing of the water content of the soil.

The largest difference between wall and ground modelling is the great depth of the ground compared to wall thickness for calculation of heat conduction. The depth has to be in the order of metres in order to obtain a constant ground temperature over the period the calculations take place. An example of calculated ground profiles are presented in figure 7.

The figure depicts some results over a period of 20 days for a rock ground. Typically is the raising temperature with the increasing better weather. The most interesting is, however, how well the model manages the diurnal changes.
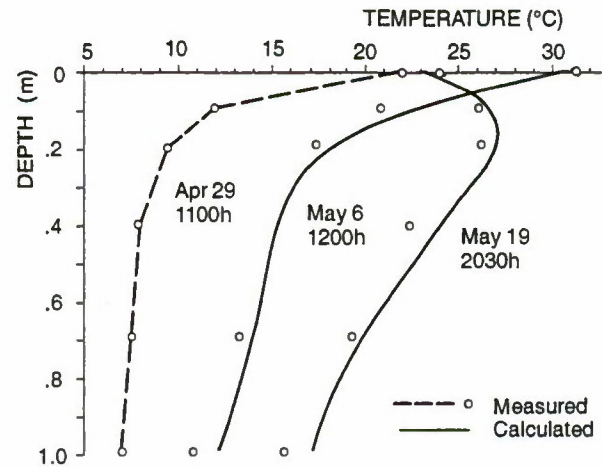
Figure 7 Rock temperature profiles

As starting parameters for the calculations was used measured rock temperatures from the first day as shown in the figure. For the further calculations it was assumed a constant temperature at a depth of 3 m. The environmental parameters were measured during the whole period, and they were used as input in the model.

Surface temperatures of other types of bare ground (sand and peat) has also been calculated. The deviation between calculated and measured temperatures are, as expected, greatest for peat in that only 57% of the deviations were less than 2°C and 28% between 2 and 3°C. For sand the deviations were 80 and 14% respectively, while for rock (figure 7) the numbers were 94 and 5%.

SYNTHETIC IMAGES

Inserting objects into backgrounds

Synthetic images are made up by inserting images of real objects into images of real backgrounds. The work on synthetic images are described in reference [4]. The foundation is a database of digitized images of different objects and of representative backgrounds. Practically all image handling has been performed on a Topaz working station from Primagraphics Ltd.

The strength of the Topaz system is that it handles different image planes as bit-maps. Thus there can be a visual and thermal background plane, a visual and thermal object plane and a form plane. The latter is necessary in order to define the form of the object.

In figure 8 is illustrated how a visual image of a military vehicle is inserted into a visual image of a background. The form plane tells which parts of the object plane that will be presented in the final image: white/seen, black/not seen.

The object can be duplicated and reproduced in different sizes before placing them at the "right distance" in the
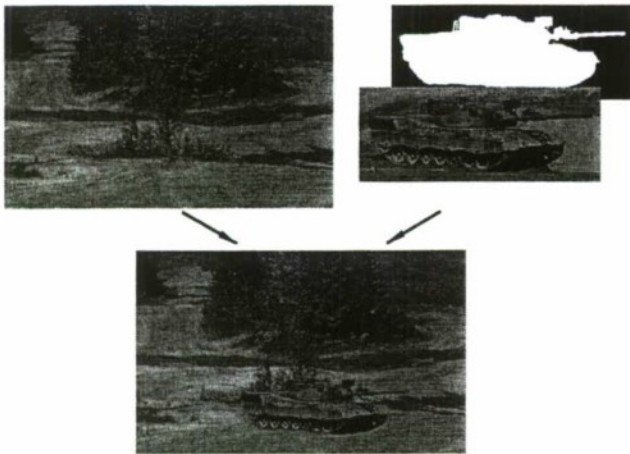
Figure 8 Inserting an object into a background

Figure 9 Thermal camouflage nets

background terrain. This possibility makes such images valuable in studying camouflage at different distances in the terrain.

The object can be a military vehicle as in figure 8, or it can as well be a tree. If a tree is chosen as an object, it will need a separate form plane. It is possible to give a tree a partly transparent quality which makes the simulation more natural. Thus it is possible to build up several trees in the terrain – also in front of the vehicles.

Contrast manipulation

Before placing the various planes in the final image it is possible to manipulate the contrasts of the original images. By modifying the contrast for parts of the object, simulation of camouflage measures like thermal screens over hot parts can be undertaken.

The background images are recorded during clear days and nights. However, atmospheric contrast reduction both in the visual and thermal images can be introduced to the final image. This is done by applying a distance plane where each pixel is given the distance to the corresponding pixel in the background plane. The Lowtran 6 model is used for calculating atmospheric transmission for the line-of-sight to each pixel in the image.

Introducing camouflage nets

Camouflage nets are introduced into the images by applying the thermal model described in the previous chapter. In the example shown in figure 9 a camouflage net is mounted over a combat vehicle.

Modelling the net is as before, but modelling the vehicle requires a modified procedure. Information on surface temperature is needed for modelling, while the only information we got about the vehicle is an image. By calibrating the thermal camera by a blackbody when recording the

image, a temperature scale is obtained, which in its turn is used to estimate the temperature in each pixel of the object image.

Based on the assumption that the body of the vehicle is made up of 40 mm thick steel armour plates, the variation in vehicle temperatures over time is calculated. The hottest parts due to heat generated by the engine are, however, not varied as the external heat exchange phenomenon are assumed negligible for these parts.

Although the shape of the net in this example is too regular, the figure illustrates the power of the described combination of image manipulation and thermal modelling.

Modelling thermal smoke

Modelling smoke clouds are based on the assumption of exponential degradation of the transmission through the smoke. Smoke grenades are modelled as depicted in figure 10. From each grenade the smoke diffuses with time from the delivery point according to the equation

$$m(r,t) = C(t) \cdot \exp\left(-R^2/2\sigma(t)^2\right) \qquad (6)$$

where

r   position relative to delivery point
m(r,t)   mass concentration in r at time t
C(t)   total aerosol concentration at time t
R   distance from cloud centre
$\sigma(t)$ diffusion coefficient at time t

The diffusional spread of the smoke cloud has been overlaid the spread by wind in that a wind profile which increases exponentially from ground and upwards.

Finally the transmission between the observer and each pixel in the image is calculated by integrating through the smoke cloud. The dynamics of an evolving smoke cloud
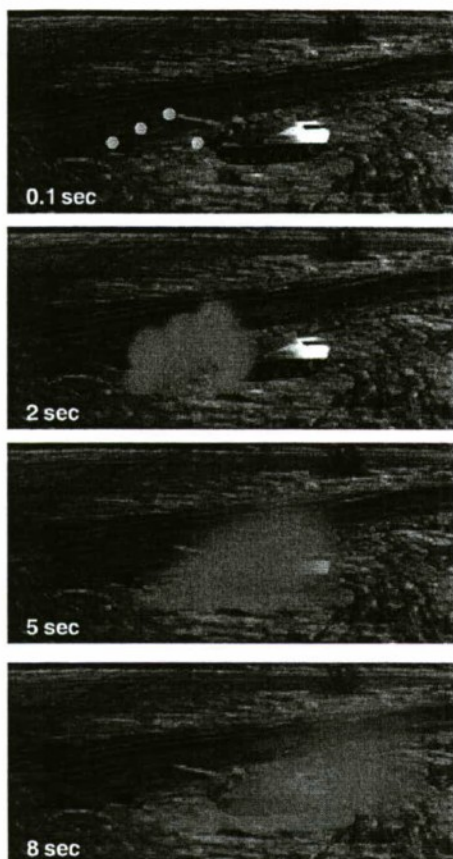
Figure 10 Thermal smoke

can best be studied when the results are presented by means of video, of which a few frames are presented in figure 10.

The model also gives the opportunity to model various types of smoke as shown in figure 11.
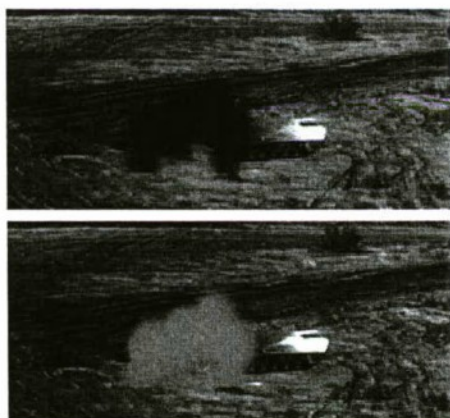
Figure 11 Cold and warm smoke

FURTHER WORK

Although the mathematical models and the system for manipulating images are powerful tools for studying camouflage, we intend to improve the tool, and the work is already in progress. The aim of the work can be defined as: Given a thermal image of a terrain recorded under certain conditions, we intend to make a model giving the appearance of the same terrain under quite different meteorological conditions as well as at a different time of day. The intention is not to construct a fully synthetic image. The work will also seek to solve the same problem for military vehicles.

The work proceeds in three parallel lines, closely intermingled:

1)   Modelling of the terrain ie mainly vegetation, which will be a challenging task, but not done in more detail than is necessary for obtaining reasonable results.

2)   Modelling of objects. Here we are prepared to acquire complete models from outside our establishment.

3)   Measurement of all relevant parameters, meteorological and others, with recording of thermal images at two measuring sites in Norway. Sampling frequency will be less than an hour.

The results from the measurements will be needed both for input parameters for the model, and for evaluating the model.

REFERENCES

[1] Rønning A T, 1982, "A One-Dimensional Model for Calculation of Surface Temperatures of Buildings" (in Norwegian), TN F–276, Norwegian Defence Research Establishment

[2] Rønning A T, 1979, "Thermal Modelling and Measurements of Selected Net-Camouflaged Objects", TN F–342, Norwegian Defence Research Establishment

[3] Rønning A T, 1982, "Thermal Radiation from Bare Ground" (in Norwegian), FFI/NOTAT–82/3020, Norwegian Defence Research Establishment

[4] Østevold E, Høimyr T, Storhaug G A, 1990, "System for generation of synthetic images" (in Norwegian), FFI/RAPPORT–90/4002, Norwegian Defence Research Establishment

Synthetic scene simulation for camouflage assessment.

Mrs M. A. Gilmore*, Dr I. R. Moorhead+, Mr D. Oxford#,

* Airborne signatures and IRCM, Weapons Systems Sector, e-mail: MAGilmore@dera.gov.uk
+ Protection & Performance Dept, Centre For Human Sciences, e-mail: I_Moorhead@dera.gov.uk
# Sensors & Avionic Systems Dept, Sensors and Processing Sector, e-mail: deoxford@dera.gov.uk
Defence Evaluation and Research Agency, Farnborough, Hants GU14 OLX

Dr D Filbee♦, C Stroud♦, G Hutchings♦, A Kirk♦

♦ Hunting Engineering Ltd (HEL), Reddings Wood, Ampthill, Bedford, MK45 2HD, UK
e-mail: drf@hunting2.demon.co.uk

ABSTRACT

Real world trials are very expensive, therefore synthetic imagery is used to assess the effectiveness of different camouflage schemes under a wide range of conditions. It is impossible to exactly reproduce the real-world synthetically, but for many applications this is not necessary. CAMEO-SIM has been developed as a broad-band (0.4 – 14 micron) physically accurate scene simulation system which models the interactions between the 'target' and the environment, such as the shadow effects of trees on helicopters and at longer wavelengths the reflection of thermal components on other surfaces. One of the key features of CAMEO-SIM is that images can be rendered to different levels of fidelity so that an appropriate balance between rendering time and accuracy can be achieved for different applications. CAMEO-SIM has undergone some verification tests and a more comprehensive validation programme is planned.

The initial requirement was to assess the effectiveness of aircraft camouflage schemes but the system could be used to study camouflage on any type of vehicle. CAMEO-SIM has recently been extended to give true colour visible band imagery in CIE X,Y,Z tristimulus values and to allow for different monitor characteristics. This paper will give a summary of CAMEO-SIM, describe the results of verification tests undertaken, present example images of different scenes in different wavebands and show the effect of different levels of image fidelity.

INTRODUCTION

All camouflage is a compromise between the requirements to match different backgrounds in different wavebands at different times of the year. The compromises made in the past were generally made by subjective assessment of the effectiveness, generally under a limited range of environmental conditions. Synthetic scene generation offers a viable alternative to field trials for the quantitative evaluation of camouflage. CAMEO-SIM has been developed as a physics based, broadband, scene simulation toolset to enable the quantitative evaluation of both current and future camouflage. The same toolset may also be used to assess concealment and deception methodologies.

OVERVIEW CAMEO-SIM

The goal of the CAMEO-SIM system is to produce synthetic, high resolution, physically accurate radiance images of target vehicles in operational scenarios, at any wavelength between 0.4 and 14 microns. The main components of the system are shown schematically in Figure1. These are described in detail elsewhere [1]. The software was developed with a scaleable rendering kernel in which imagery can be produced at different fidelities and frame rates depending on the image application and wavelength of operation.

At longer wavelengths, such as the mid infrared, the full hemispherical integration of the incident irradiance enables the software to account for the thermal radiative interaction between different surfaces. An example of this interaction, in which hot surfaces are reflected in the road is shown in Figure 3.
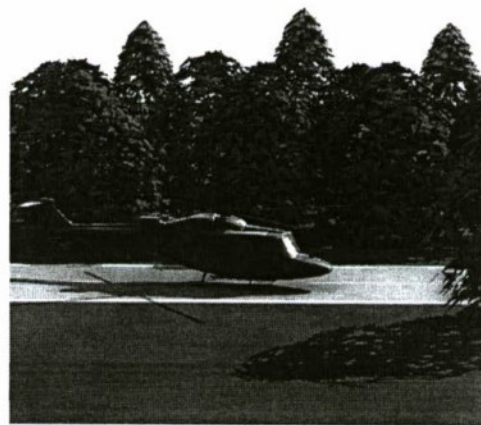
Figure 2: Visible band, clear day, BRDF properties applied to airframe, shadows from sun and sky.

Figure 1. Block diagram of CAMEO-SIM components illustrating the data flow through the different processes.

A key element of CAMEO-SIM is that there is a complete audit trail between the data used to form the image and the final rendered image. This means that it is possible to conduct carefully controlled manipulations of scene properties which can be repeated.

The background scene is generated using MultiGen®. All geometric objects forming the synthetic environment are modelled using textured faceted structures. Texel values in these textures are mapped to real materials which have measured physical properties associated with them e.g. bi-directional reflectance. Each texel is then considered as a mixture of up to three different materials. This means that the physical properties of both soil and grass are modelled when using a 'grass' texture. Complex objects are modelled as a number of polygons. This means that the three dimensional effects of trees, including shadow effects can be simulated (Figure 2).

Figure 3. CAMEO-SIM 3-5 micron image showing thermal reflections, no sensor effects.

IMAGE DISPLAY

CAMEO-SIM computes the radiance in specified sub-bands for each pixel in the image. These are then summed to produce an 'in-band' grey-scale image.

The software can also display each sub-band as a grey scale image or any three sub-bands as a false colour image. CAMEO-SIM can also display visible band true-colour imagery. This is done by first evaluating the spectral radiance image cube for a user-defined number of sub-bands between 0.38 and 0.78 microns. The spectral image cube is then converted into device independent colour space represented by the tristimulus values X,Y,Z, using the spectral tri-stimulus values of the CIE 1931 Colorimetric Standard Observer. The monitor that is used for the image display is calibrated so that the X,Y,Z values can be converted to R,G,B values. Various luminance transforms can be employed to make best use of the limited CRT dynamic range (Figure 4).

Figure 5: Visible band, clear day, all surfaces diffuse, no shadows, no sensor effects.

Figure 4. Visible band image, no sensor effects.

IMAGE FIDELITY

The lowest fidelity mode is real-time and can be used to develop and preview the scene before it is passed to a scaleable higher fidelity rendering kernel. This high fidelity renderer can model directional reflectance effects, geometric occlusion of point and extended sources, and the spectral integration of the optical properties with the atmosphere. The level of fidelity required can be selected before rendering the image, so that appropriate imagery can be generated for different applications (Figures 5 and 6).

Figure 6. Visible band image, clear day, all surfaces diffuse, shadow effect from the sun, no sensor effects

205

SENSOR EFFECTS

At present CAMEO-SIM is used to create an image from a static sensor, with either static or moving 'targets'. The images can be displayed on a monitor for observer trials, or the digital image (in units of $Wm^{-2}sr^{-1}$) can be analysed. The ongoing development programme is linking CAMEO-SIM with other programs so that a moving sensor can be modelled. In addition sensor effects such as MTF, sampling, jitter and noise can be added to the images (Figures 7 and 8).

Figure 7. CAMEO-SIM 3-5 micron image, no sensor effects.

Figure 8. Same image as that in Figure 7 with MTF, sampling, jitter and noise sensor effects added.

VERIFICATION AND VALIDATION

A number of analytical tests have been conducted to verify the operation of the high fidelity rendering equations. These are reported elsewhere [2] and include:

1. Blackbody radiance tests
2. Contrast in an isothermal environment
3. Shadowing and blocking
4. Spectral calculations
5. Radiometric calculation of lighting effects
6. Directional emission of uniformly textured and heated spheres
7. Material assignments on a texture
8. Bi-directional reflectivity of uniformly textured and heated spheres
9. Small target rendering

These verification tests will continue and be extended to include validation tests. The validation processes will involve three separate approaches:

1. Highly simplified scenarios that can be synthesised within CAMEO-SIM and measured (Figure 9)
2. Comparison of the statistics of real and synthetic imagery. A range of techniques will be used including models based on human vision performance and higher order statistics. The significance of different effects, such as shadows, in different wavebands will be determined.
3. Observer performance – such as comparison of detection ranges of aircraft in different camouflage schemes
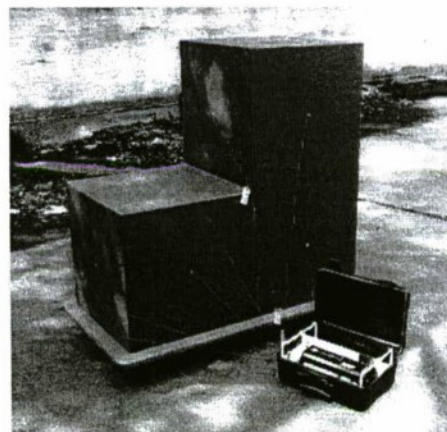
Figure 9. Photograph of the step object to be used in validation trials. The photograph shows how the object can be instrumented with thermocouples for temperature measurement. (Lighter patches in the image are shadows)

CONCLUSIONS

CAMEO-SIM can generate physically accurate imagery between 0.4 and 14 microns to different levels of fidelity, to allow a trade-off between accuracy and rendering time. The software has undergone a range of verification tests to show that the correct values are computed. A programme of validation is underway to ensure that meaningful results are obtained using the software tools. This programme will address three different aspects of the synthetic imagery – statistics, real-world comparisons (including scenes with simple objects) and performance prediction. The functionality of CAMEO-SIM is being extended as part of an ongoing research programme. For example recent improvements include true-colour visible band image display, sensor effects and a moving sensor.

ACKNOWLEDGEMENTS

REFERENCES

1. Oxford D. et al. 'CAMEO-SIM: A physically accurate broadband EO scene generation system for the assessment of air vehicle camouflage schemes.' *Proceedings of the Ninth Annual Ground Target Modelling and Validation Conference 1998.* Houghton, MI
2. Gilmore M A et al. 'CAMEO-SIM: a broad-band scene generation system that is 'fit for purpose'. *Targets and Backgrounds: Characterisation and Representation V. Proceedings SPIE 3375. 1999*

MuSES: A New Heat and Signature Management Design Tool for Virtual Prototyping (a follow-on)

Keith Johnson, Allen Curran, David Less, Derrick Levanen, and Eric Marttila
ThermoAnalytics, Inc.
Calumet, MI 49913
and
Teresa Gonda and Jack Jones
US Army TACOM
Warren, MI 48397

ABSTRACT

This is a follow-on paper regarding last year's introduction of the development of a new thermal design and signature modeling tool. The U.S. Army awarded a Small Business Innovation Research (SBIR) Phase II program at the end of 1997 to create a new design tool called MuSES (Multi-Service Electro-optics Signature) code capable of meeting the arising requirements for rapid prototyping as well as addressing the weaknesses of current thermal signature models. Because of the strong connections between signature management for military ground vehicles and heat management for commercial automobiles, there is considerable dual-use collaboration and commercial commitments from the automobile industry. The first version of this program is currently available commercially. This paper will discuss what has been accomplished and what has yet to be done this year. A key ingredient to the success of this modeling tool, particularly for rapid prototyping, is the conversion of multiple solid CAD geometry into a faceted mesh that is clean (i.e., does not require additional human interaction and does not add small sliver type polygons).

TAI THERMAL PRODUCTS

In the past, ThermoAnalytics supported many different sponsors with individual codes: TACOM-PRISM, AFRL-TCM2, Ford-RadTherm, Chrysler-CTD, Commercial-WinTherm). By building a new WinTherm that serves as the thermal solver kernel for all of these programs plus future ones, upgrades are automatically passed on to all specialized modules.
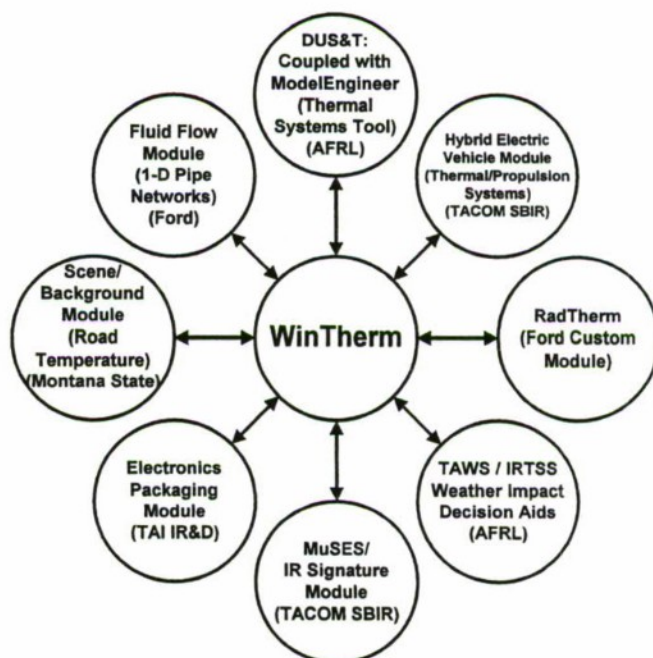
Figure 1. ThermoAnalytics' Complementary Programs Provide a Leveraged Advantage for MuSES.

Description of WinTherm/MuSES Features
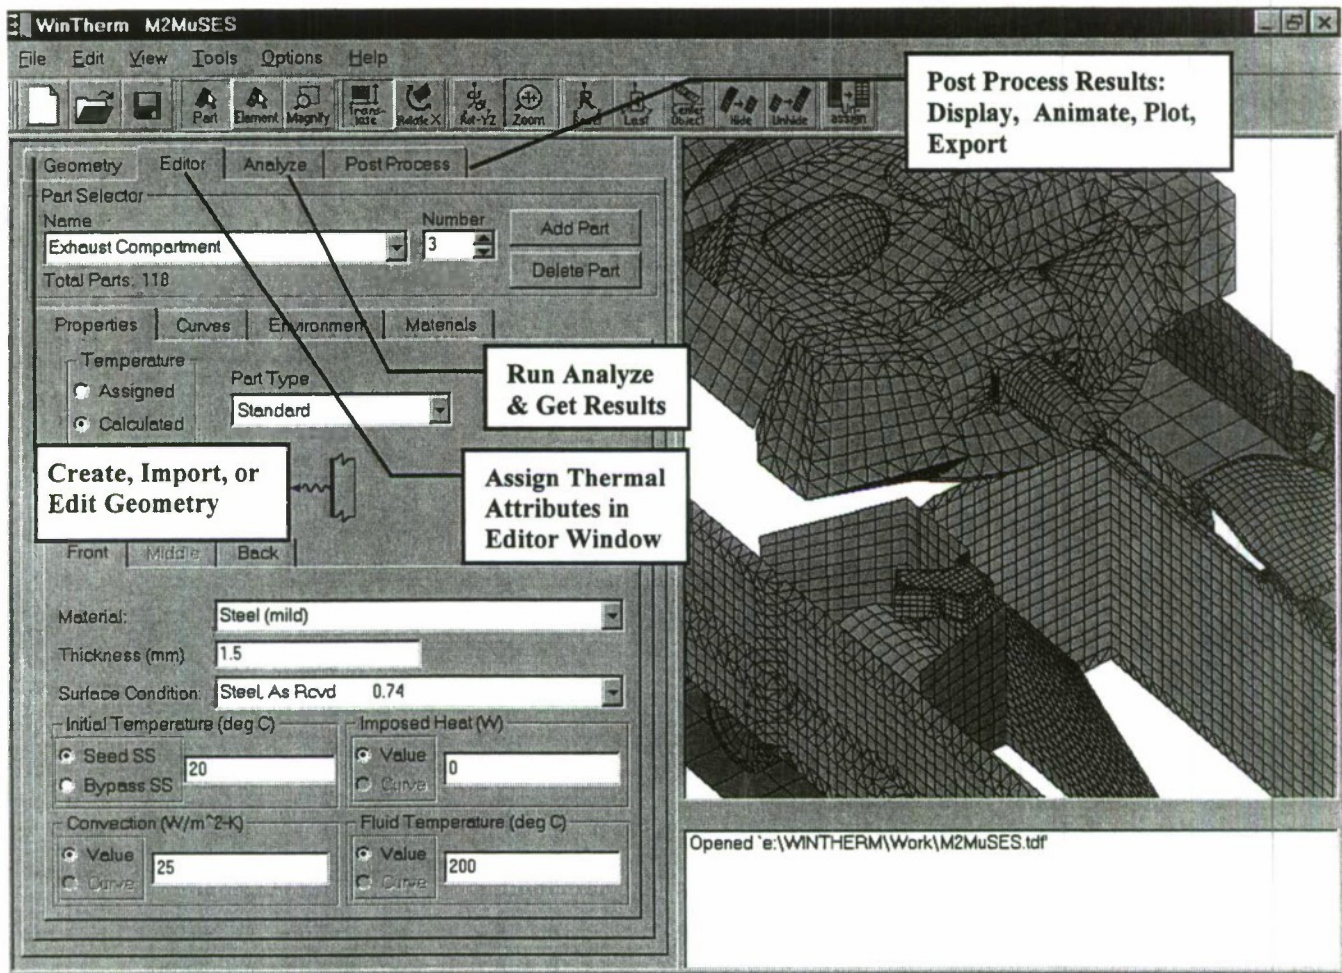
Figure 2. Features of WinTherm / MuSES

Just a few of the features can be seen above from the graphical user interface. Many more will be presented next.

GUI: Intuitive User Interface that Combines ALL Aspects of Modeling into Single Program

- Geometry
- Model Construction / Editing
- Analysis Pre-Processing and Solution
- Post Processing of Results

The Graphical User Interface is one of the key features that WinTherm is so easy to use. We obtained a lot of engineering user feedback during the development of the GUI that contributed to the layout and user interactions.

Code: Cross-Platform and Complete Rewrite in C/C++

Because there are both UNIX and Windows users, we felt it was important to develop only "one" code to span across all platforms. This was accomplished by using a commercial cross platform tool called Qt. A lot of time went into the selection of this tool and we are very satisfied with the choice.

Hierarchical Data Format: Machine/Language Independent File

The HDF file format provides many benefits: cross-platform and language independence, extensible (can add new data to the format) without affecting previous readers, compact (binary).

Thermal Solver: New Modular FDE Solver that Features Partially Direct Solution and Net- Radiation Method

■ *Partially-Direct Method*

- Computationally Expedient Modeling of Thin or Thermally Diffuse Elements

■ *Net-Radiation Method*

- Accurate Modeling of Low Emissivity Cavities
- Allows Temperature Dependent Emissivity (see Figure 7)

■ *Modular and General Purpose*

- Add "Hook" Functions - User Defined Boundary Conditions.

■ *Conduction: Automatic 3-D Linkages with Anisotropic Conductivity in the Works*

■ *Convection: Enhanced Wind Convection and Internal Convection Libraries*

Background Models: Spatial Backgrounds using Digital Elevation Model (DEM) inputs

Includes All PRISM 1st Principles Terrains

- (Soil, Foliage, Concrete, Asphalt, Snow, Water, etc.)

Default Background

- Influences (but Not Influenced By) Target
- Can Be Used in Conjunction with Faceted Terrains (Use Small Faceted Terrain Local to Target)

Background Part Types (Geometry Optional)

■ *No Geometry*

- No Interaction With Target/Other Backgrounds
- No Need to Make Additional MuSES Runs

■ *With Geometry*

- Full Interaction With Target/Other Backgrounds
- Can Exclude Rigorous View Factor / Apparent Area Calculation on a Part Basis (e.g., for sections distant from target)
- Geometry Can Be Read From Digital Elevation Maps (such as produced by ARC/INFO, etc.)

View Factors: Accelerated Ray-tracing Technique using 3D Spatial Subdivision of Voxels

The acceleration technique is 3D spatial subdivision

1. Spatially divide object/scene into voxels
2. Only voxels pierced by ray being cast are searched for object intersections
3. Greatly reduces number of ray/object intersections that must be evaluated

Speed Comparison

■ *BRDM2 Test Case - 3,743 patches (1 elem / patch)*

- Indigo2 195Mhz R10000, Solid Impact Graphics
  1. Single Plane Hemicube (128x128 window): 105.7s
  2. *Equivalent Accuracy (Fast) Ray Trace (1.9m rays): 63.5s*
  3. Standard Accuracy Ray Trace (7.7m rays): 189.5s
- Standard PC Hardware (PIII 500)
  1. *Equivalent Accuracy (Fast) Ray Trace (1.9m rays): 28.5s*
  2. Standard Accuracy Ray Trace (7.7m rays): 101.7s
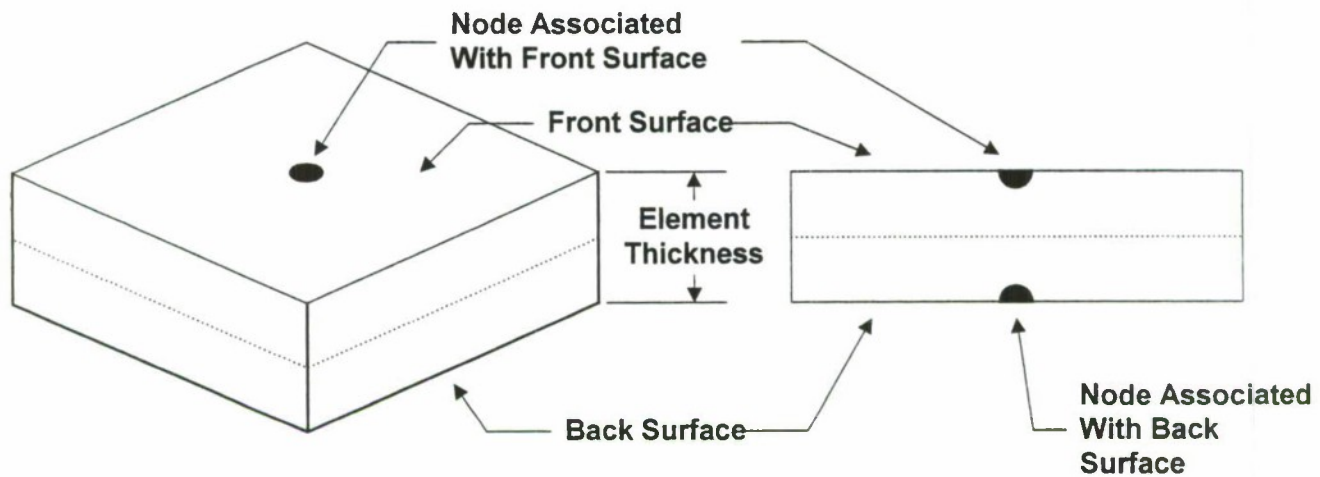
Figure 3. Surface of a 3D Element is What We Commonly Refer to as a Facet
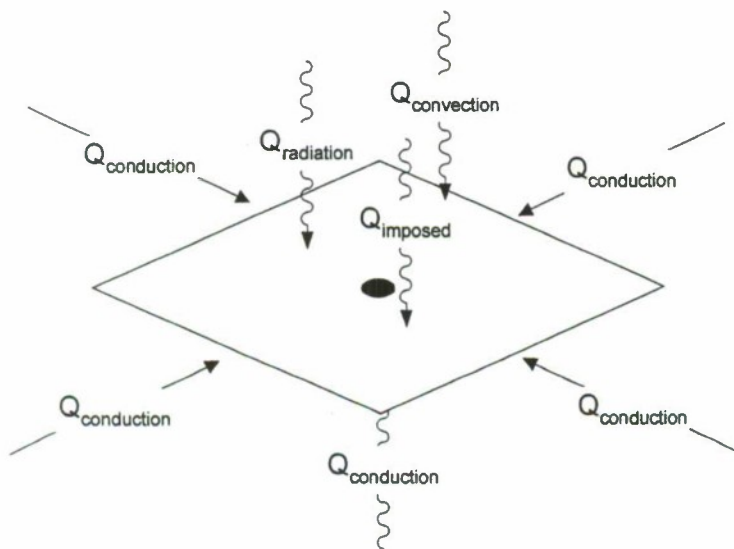
Governing Heat Transfer Equation

$$m_k C_{p_k} \frac{\partial T_k}{\partial t} = h_k A_k (T_f - T_k) +$$

$$\sum_{j=1}^{Ncond} \left(-k_{kj} A_{kj} \frac{T_k - T_j}{L_{kj}} \right) +$$

$$Q_k + Q_{imp_k}$$

Enclosure Method Radiation Equation

$$\sum_{j=1}^{N} \left(\frac{\delta_{kj}}{\varepsilon_j} - F_{k-j} \frac{1-\varepsilon_j}{\varepsilon_j} \right) \frac{Q_j}{A_j} = \sum_{j=1}^{N} \sigma F_{k-j} (T_k^4 - T_j^4)$$

Figure 4. The Element (and its Associated Thermal Nodes) is the Building Block for the 3D Heat Transfer Solution that Includes a Sophisticated Enclosure Method Radiation Solver

- **Only the surface mesh is required for geometric modeling. Internal nodes are automatically applied as needed with specialized Part Types.**
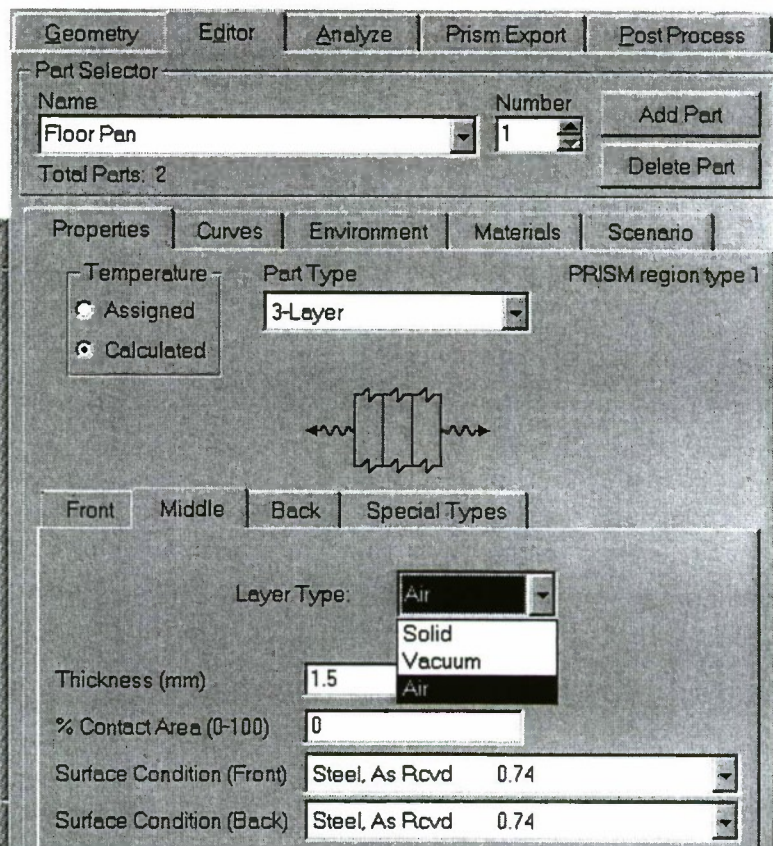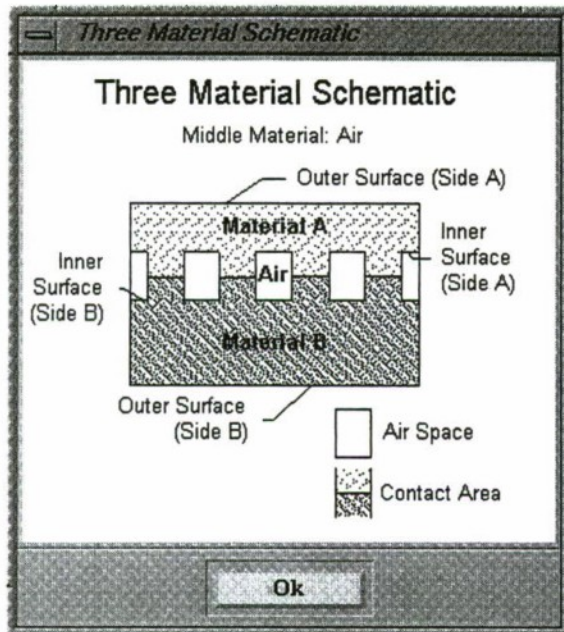
Figure 5. Composite (Multi-Layer) Part Inputs are Easily Added
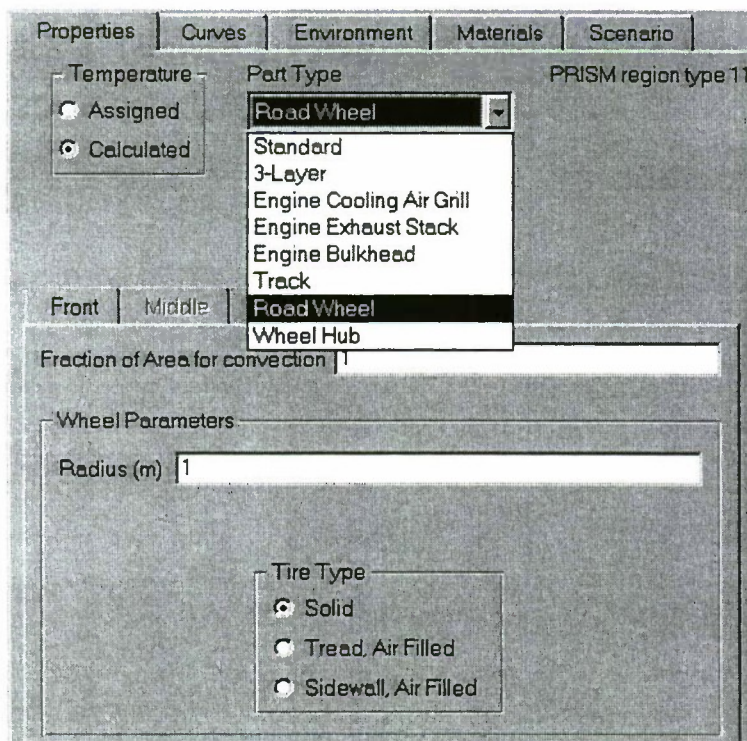
Other Specialized Part Types include the PRISM types

- Equivalent PRISM facet types displayed
- Includes parameters (e.g., radius) as appropriate
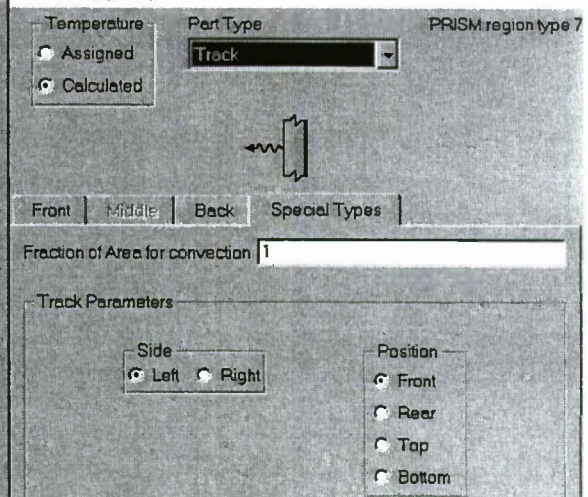
Figure 6. Specialized Part Types Include the Familiar PRISM "Facet Types"

The chart legend:
- Plate 1 (Constant Emiss)
- Plate 2 (Constant Emiss)
- Plate 1 (Variable Emiss)
- Plate 2 (Variable Emissivity)
- Plate 3

Y-axis: Temperature (Kelvin)
X-axis: Time (min)

Diagram labels:
- 2 mm
- Insulated
- Plate 1 (Aluminum)
- Emiss = 0.80
- h = 10 W/m²-K, T_f = 300 K
- Plate 2 (Oxidized Aluminum)
- Emiss = 0.80
- h = 10 W/m²-K, T_f = 300 K
- Plate 3 (Aluminum)
- 500 mm
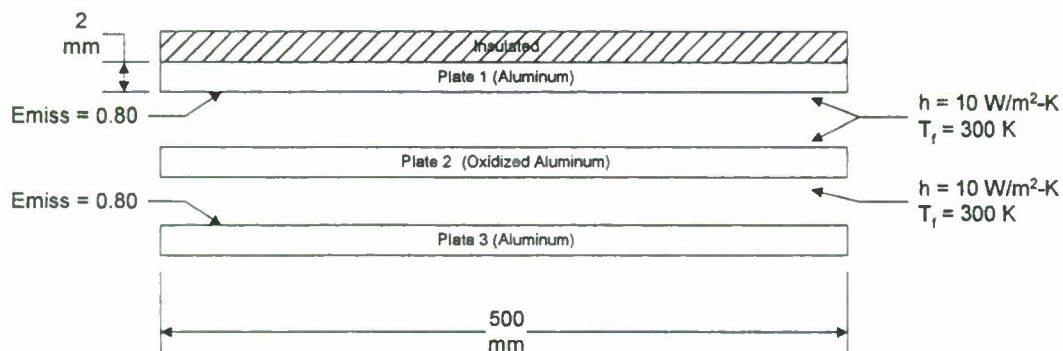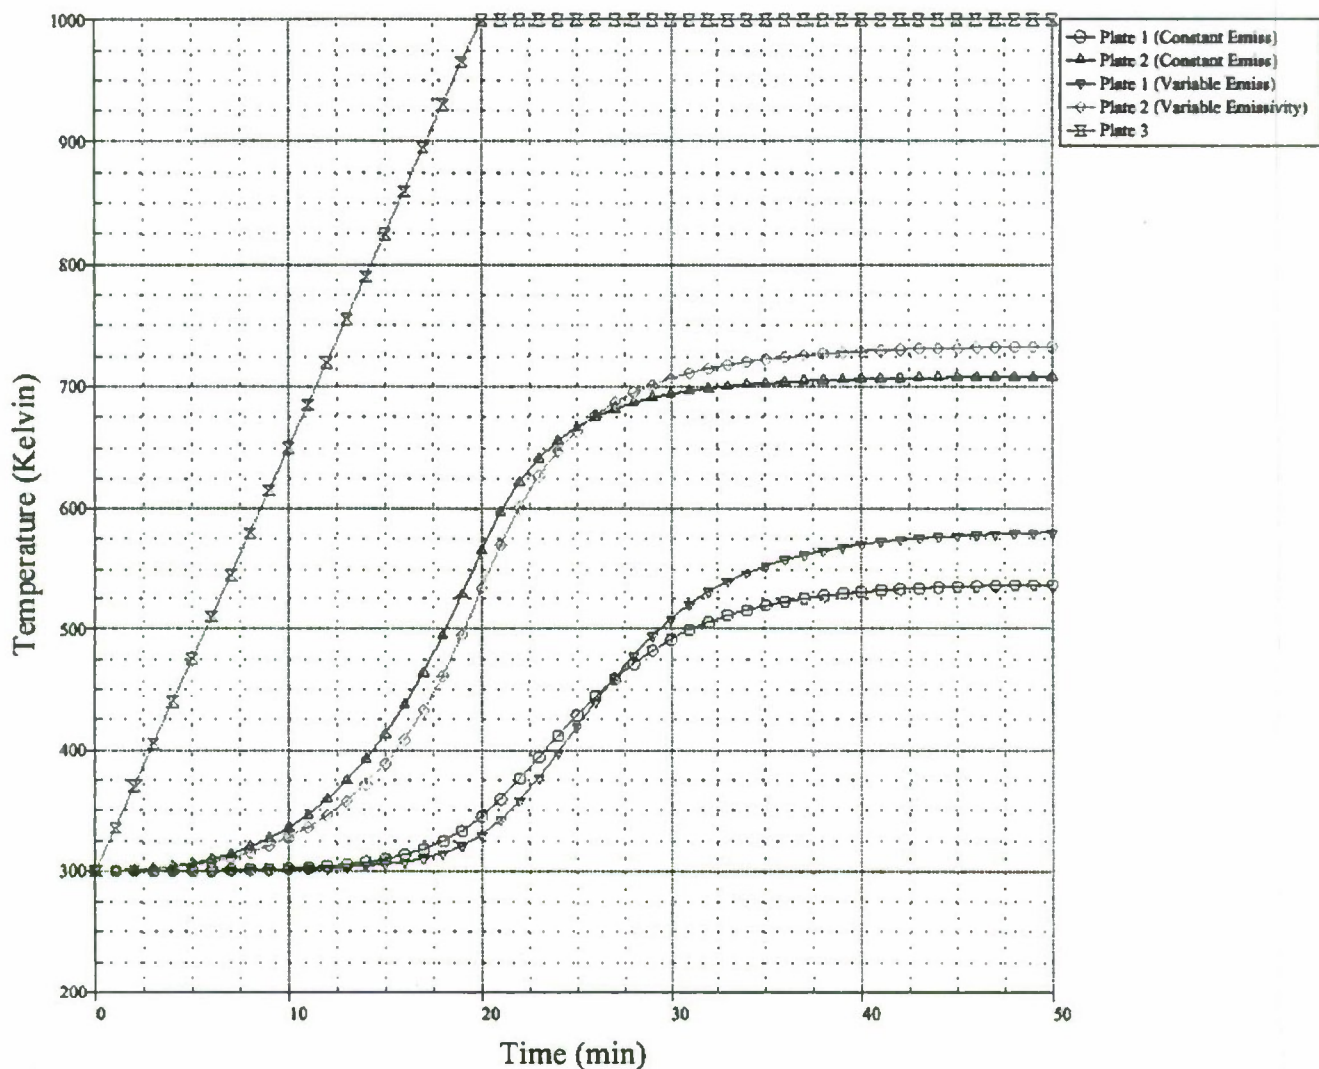
- The Transient Temperatures of Plates 1 & 2 Increase when the Actual Temperature Dependent Emissivity Properties are Used in the Analysis (in place of Constant Properties)
- Emissivity of Oxidized Aluminum varies from .2 to .32 over temperatures from 100 C to 500 C

Figure 7. Temperature Dependent Emissivity for a Multiple Bounce Radiosity Solution is Allowed with the Net-Radiation Method of Solution

Original Fine Mesh (right) is "Grouped" into Radiation "Patches" (below) based on Patch Parameters - Producing Efficient View Factor Storage
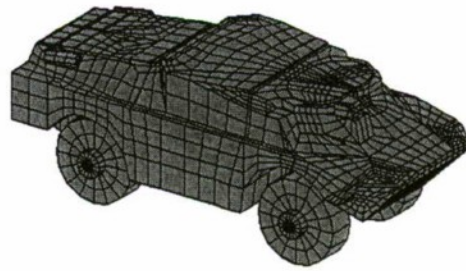
Figure 8. Patch Generator: Efficient Solution for View Factor Compression of Large Models

CFD Interface: Direct Link to CFD programs such as FLUENT for Exchanging Data: h, T_f, T_{wall}

- Import Convection Boundary Conditions obtained from CFD Analysis
- Update Convection Boundary Conditions in an Existing Model when New CFD Results are Available
- Batch Mode of Operation Supports Automated CFD Import

Figure 9. The Next Version of WinTherm will Allow a Direct Interface and Exchange of CFD Data (Film Temperatures and Heat Transfer Coefficients)

- **1-D Internal Fluid Flow is Created from Geometry**
- **User Inputs Parameters for Various Fluid Elements (pipe, pump, reservoir, etc.)**
- **Flow and Temperature are solved Simultaneously**
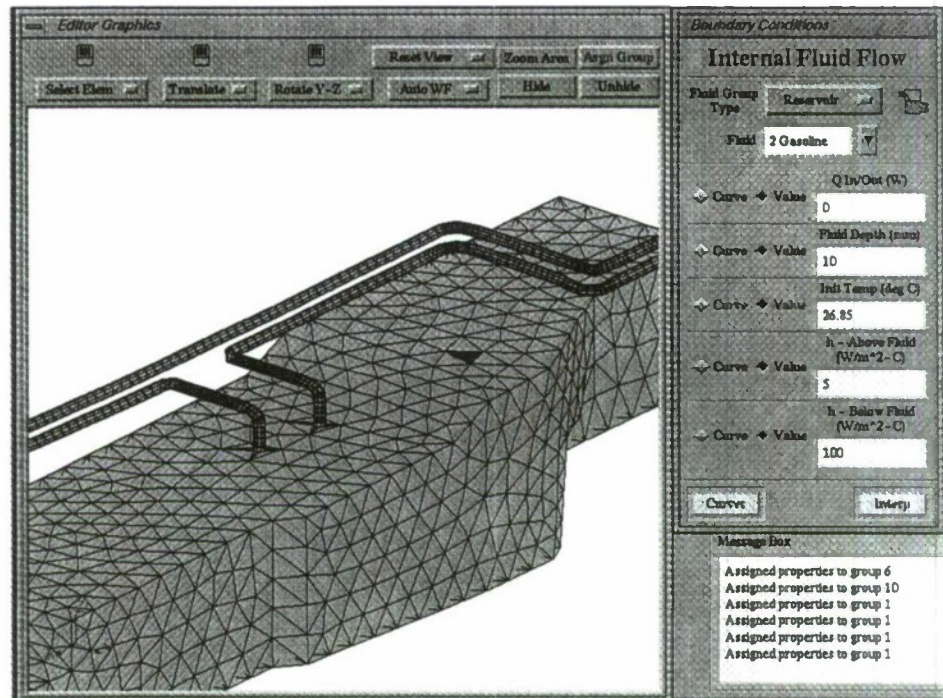
Figure 10. Internal Fluid Flow is Coming in WinTherm 5.0.

TACOM is Developing Complementary Geometry Conversion and Meshing Tools that will directly Interface with MuSES in the Future
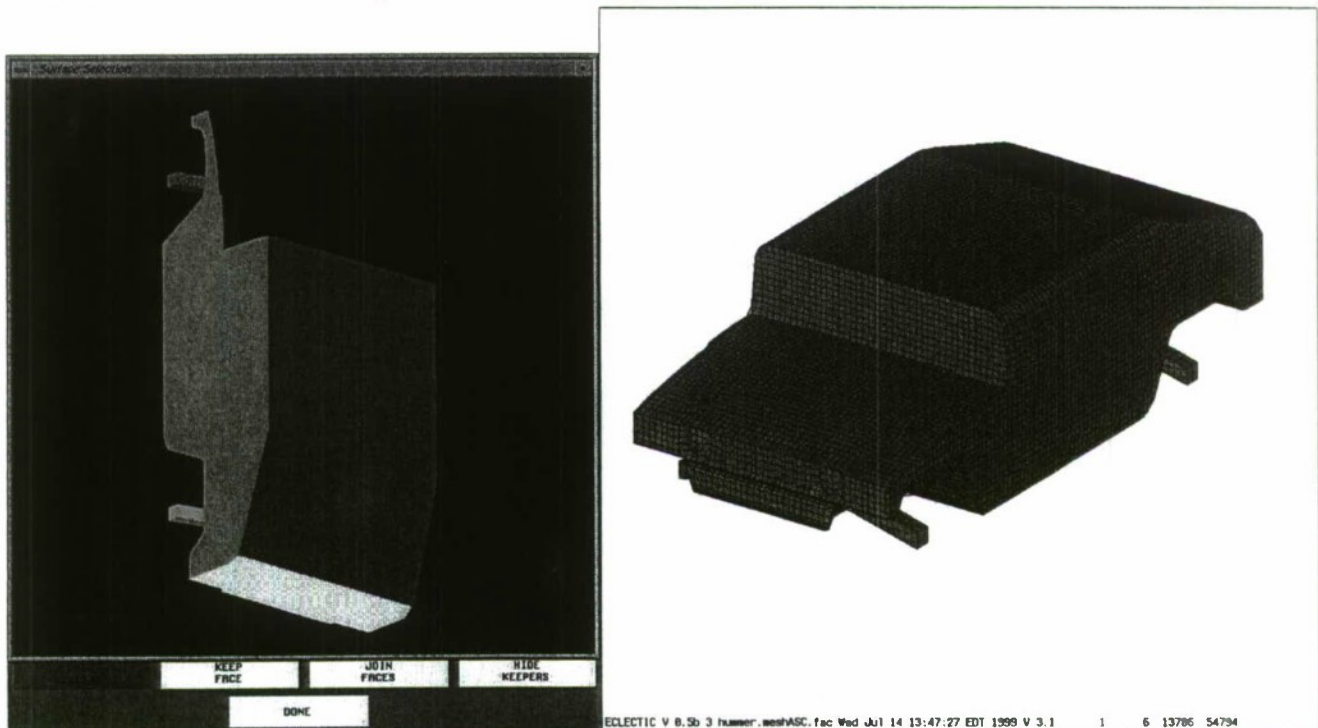
Figure 11. Goal of Eclectic is Transformation of CAD Geometry into Ideal Facetization

Validation of Thermal Solver

- Comparison of Experimental & Computational Results for Several Points
- The Model vs. Measured Temperatures vary by 1% to 10%
- Data Courtesy of Ford Motor Company
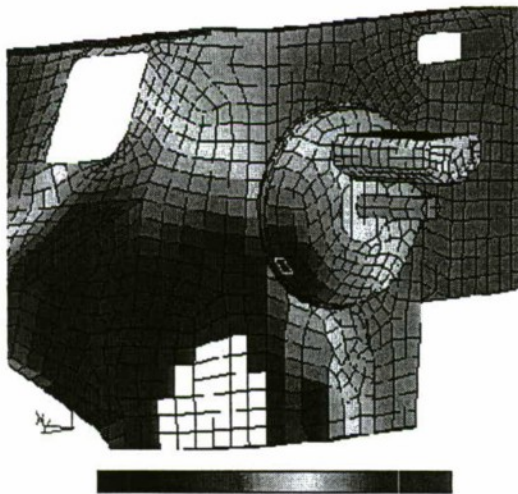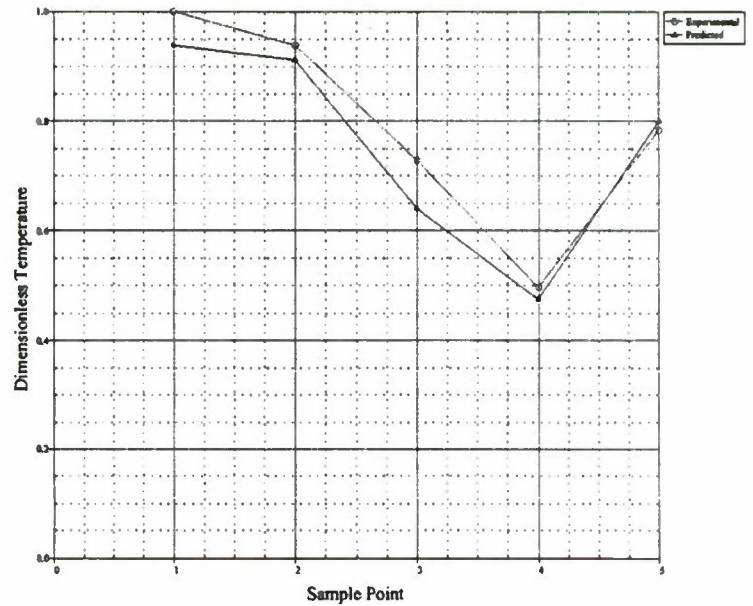
Steady State Temperature Distribution of Fire-Wall and Brake Master Cylinder
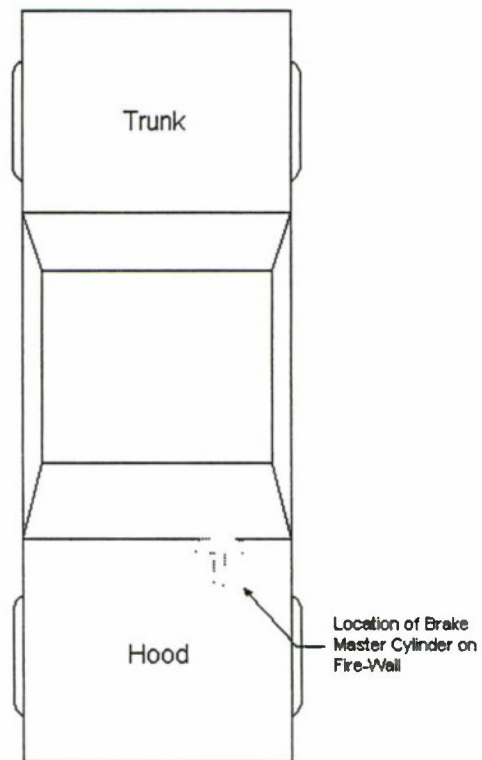
Figure 12.   Comparison with Experimental Results

Sample MuSES IR Models
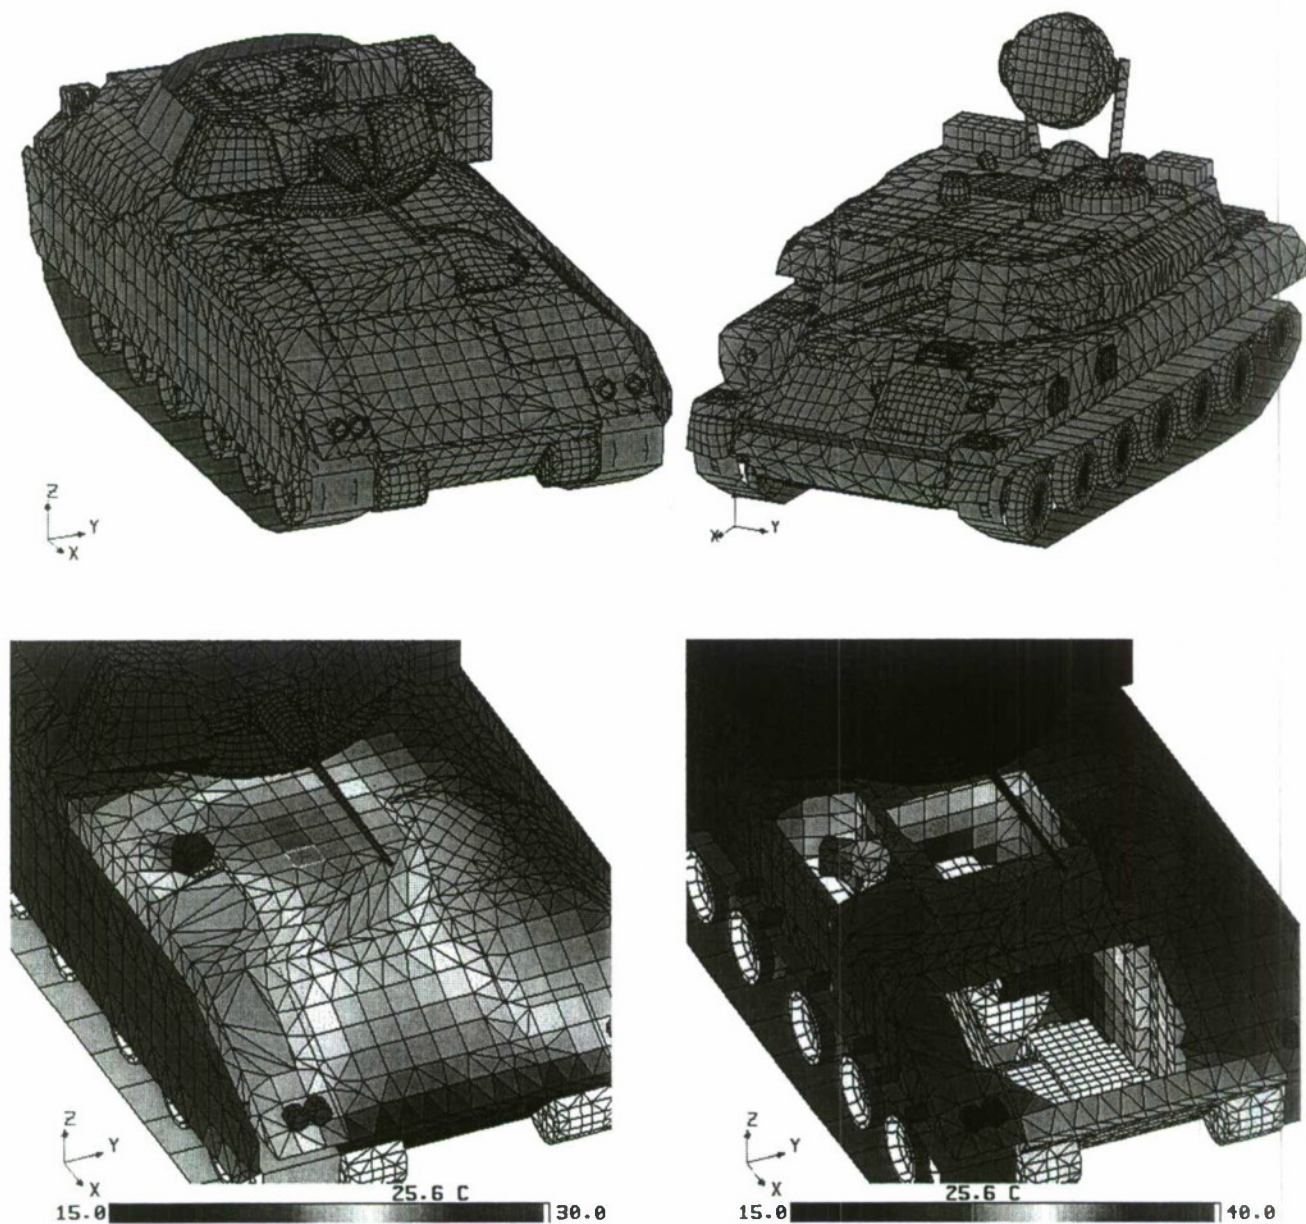
Figure 13. Modeling with MuSES

For further information about MuSES, PRISM, WinTherm, or ThermoAnalytics, visit the web site at http://www.thermoanalytics.com/muses/.

WinTherm can be downloaded from this site for demonstration and evaluation. If there are further questions, email us at info@ThermoAnalytics.com or phone us at 1-800-214-3766.

Infrared Signature Estimates of a High-Temperature Exhaust Plume

Jeffrey A. Hoffman
William R. Reynolds
Signature Research, Inc.
Calumet, MI. 49913

ABSTRACT

Estimating thermal signatures of a high-temperature, multi-component gaseous plume can be a formidable task. Issues include assessing the plume's structure resulting from an endless variety of exhaust exit geometry and ambient flow fields. In addition to fluid mechanics, the emission and absorption properties of exhaust gases are a function of concentration and local temperature, which vary throughout the plume. This paper summarizes a modeling effort that addresses both the fluid mechanics and gas emission issues. This approach uses commercially available software to generate a solution for the fluid mechanics. A custom written postprocessor discretizes the software's solution into small control volumes and estimates the absorption and emission from each volume. Lastly, apparent temperatures calculated along a user defined line of sight provided contours of apparent temperature. Three example signatures in the bands of 3-5 microns and 8-12 microns are provided for three distinct plumes.

INTRODUCTION

Estimating thermal signatures prior to product fabrication can be a valuable step towards a successfully developed product. A variety of existing software packages allow designers to estimate a vehicle's thermal signature spectrally prior to fabrication [1]. However, modeling efforts to date often ignore signatures associated with vehicle exhaust plumes. In addition, software products that implement plume models often rely on specialized solutions applicable to a narrow range of exhaust geometries and flow configurations [2]. Reasons for these shortcomings and/or over-simplifications are the complex emission properties of gases coupled with the challenge of fluid mechanics issues. This paper describes a recent approach to modeling thermal signatures of gaseous plumes that addresses both the fluid mechanics and gas phase emission issues. This approach incorporates a computational fluid dynamic (CFD) model to generate temperature and species distributions throughout the plume. With additional post processing the local absorption and emission properties are estimated spectrally using MODTRAN [3]. The emissive power and absorption properties are integrated along a user-defined line of sight, providing contours of apparent temperature on a spectral basis.

DISCUSSION

Figure 1 contains a process flow diagram of the steps used to model the thermal signature of a gaseous jet. Modeling process inputs include the exhaust geometry, spectral range of interest, engine operating conditions, and ambient conditions. The engine operating conditions include adequate information to estimate the exhaust exit temperature, gas species concentrations, and volumetric flow rate. A custom written post processor used the steady state temperature and species information obtained from the CFD model as inputs. The post processor discretized the CFD solution into small cube shaped control volumes. MODTRAN, a software capable of estimating absorption and emission properties of gases on a spectral basis, was applied to the contents in each of the control volumes. Using the absorption and emissive power properties from each cell along a user-defined line of sight, the transmitted emissive power was determined. Lastly, using a blackbody assumption for the spectral band of interest, apparent temperatures along each line of sight were calculated.

Figure 2 schematically demonstrates the layout of the discretization process. The plume lies between a hypothetical background plane and the image plane. The background's emissions pass along the line of sight where it is attenuated and/or increased depending on the absorption and emissive properties of the control volume.

Modeling Process Inputs

eg.
Exhaust Geometry
Spectral Range
Engine Operating Conditions
Ambient Conditions

CFD Modeling

Commercial CFD Software

Post Processing

1) Descretize CFD solution
2) Apply MODTRAN to each cell
3) Integrate along line of sight
4) Calculate apparent temperature
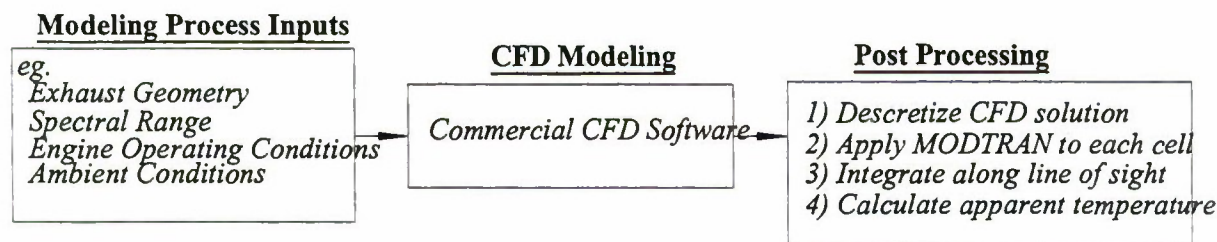
Figure 1: Process flow used to model the thermal signature of an exhaust plume.

Note that for illustration purposes, only one line of sight is shown. For the example signatures presented in this paper, there were 169 separate lines of sight (an array 13 by 13 lines of sight) used to create a single image of apparent temperature contours.
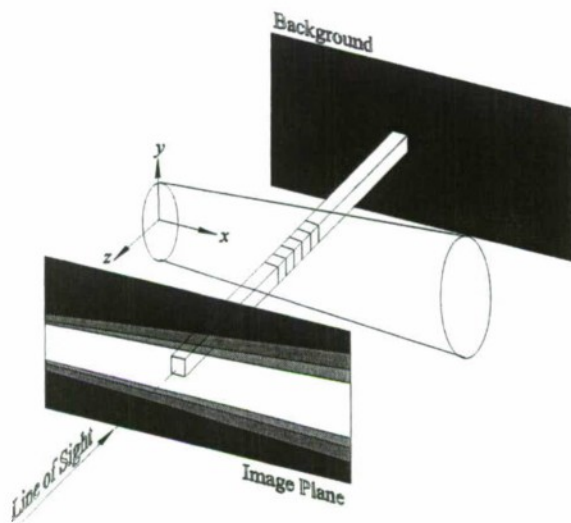
Figure 2: Schematic describing the discretization and post processing of the CFD results used to generate apparent temperature contours of the exhaust plume. (Note for illustration purposes, only one line of sight is shown).

EXAMPLES

In this section, three example thermal signature estimates are provided. For all three cases, the exhaust exit velocity was 10 m/sec, the exhaust exit temperature
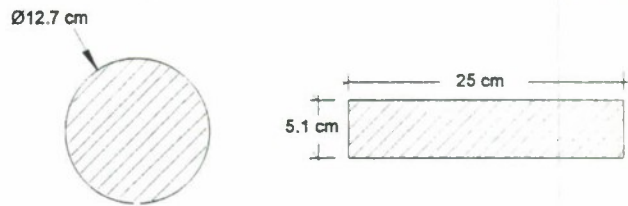
Figure 3: A schematic of the circular exit used for Cases I and II and the rectangular exit of Case III.

was 700 (°K). Exhaust species included in the analysis were H_2O, CO_2, and N_2 at molar percents of 14, 7, and 75 respectively. Initial concentration of soot particles was 0.071 g/cm^3 and modeled as spherical shaped particulates with a diameter of 2 microns behaving as blackbodies. Cases I and II utilized a circular shaped inlet while Case III used a rectangular shaped exit with the dimensions shown in Figure 3. To maintain the same volumetric flow rate between all three cases, the exit areas were kept constant at 127 cm^2. For the CFD model, the thermodynamic properties of both the ambient exhaust gases were assumed to be nitrogen.
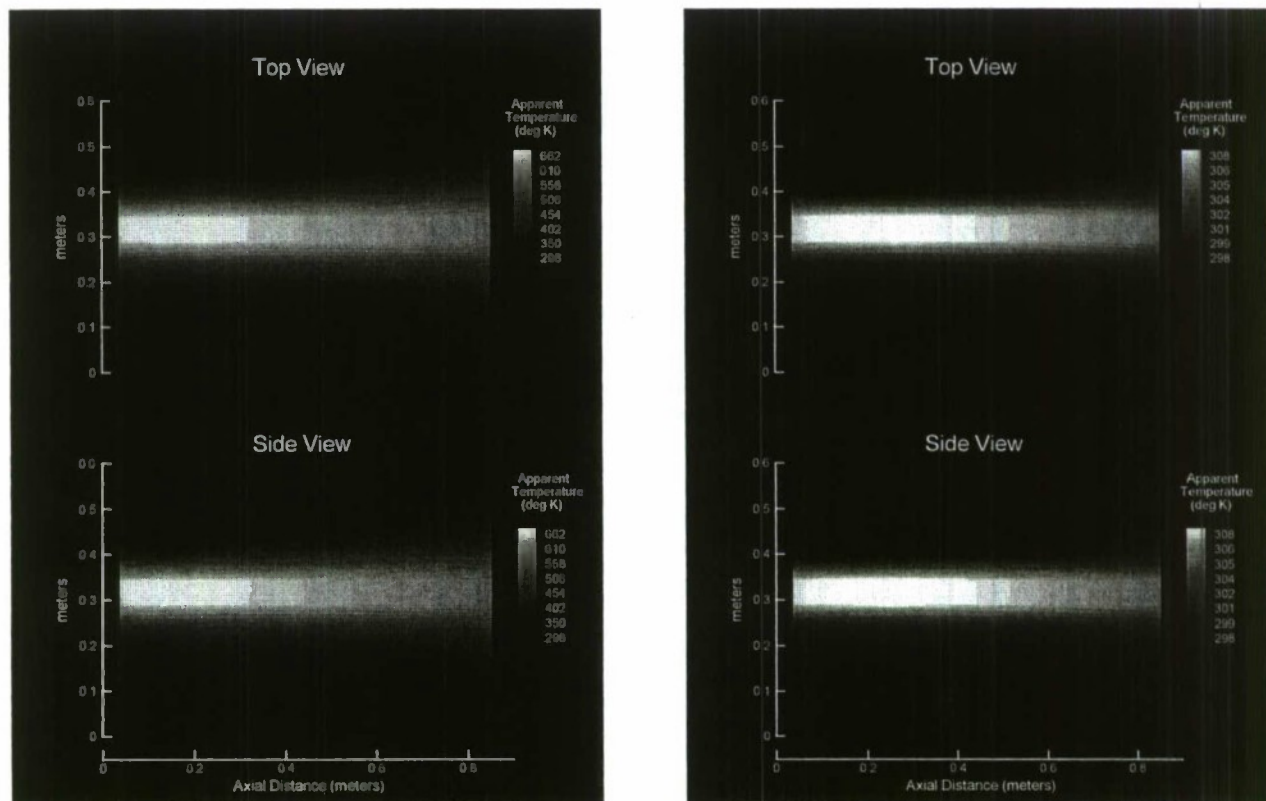
Figure 4: Contours of apparent temperatures of a circular cross section exhaust plume at a spectral range of 3-5 micron(left) and 8-12 micron (right).

219

Figure 4 represents the apparent temperatures observed from the top and side views of the plume for the spectral ranges of 3-5 microns (left) and 8-12 microns (right). As expected, the highest temperatures are noted along the plume's centerline. Due to the symmetry of the model, the top and side views are identical. Peak apparent temperatures for the 3-5 and 8-12 micron ranges were 610 °K and 310 °K respectively.

Figure 5 depicts apparent temperatures for case

due to the thin plume structure as viewed through the plume in the Top View. The shorter distance reduces the amount of hot gaseous emitting along that line of sight. In contrast, the side view, which has a line of sight through more high temperature species, exhibits peak apparent temperatures of Cases I and II. This drop in apparent temperature is due to effects of a flattened cross section, which reduces the viewing distance through hot gases. In addition, the rectangular cross section enhances
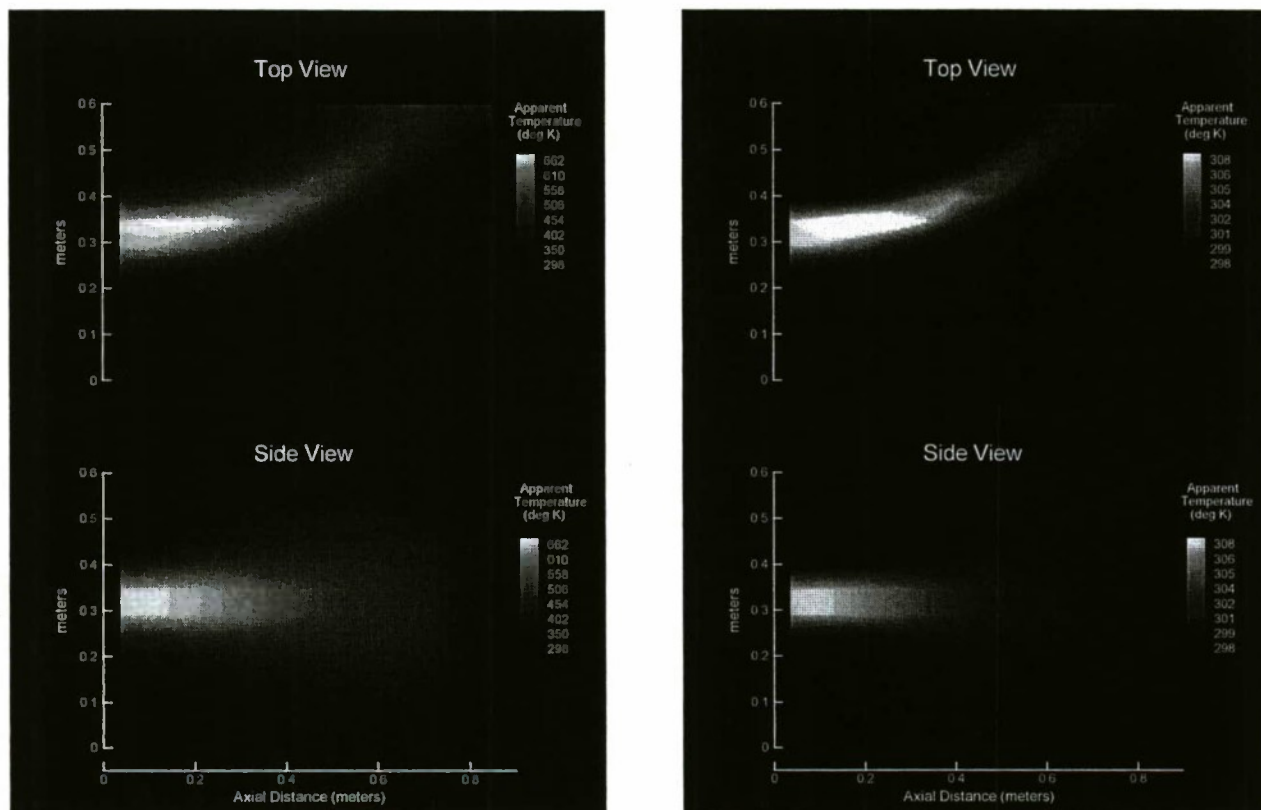
Figure 5: Contours of apparent temperature of a circular cross section exhaust with a 2.0 m/sec cross wind. Spectral ranges 3-5 micron (left) and 8-12 micron (right). Note the crosswind inlet was placed at the bottom of the Top View).

II. In this case, the boundary conditions were identical to those of Case I with the only exception of adding an ambient cross flow field of 2 m/sec. Peak apparent temperatures were the same as the previous case. However, interactions between the plume and ambient flow field generate an entirely different plume structure as it interacts with the ambient flow field.

Lastly, Figure 6 provides the estimated apparent temperature for Case III. Peak temperatures as viewed in the top view were 100 °K and 10 °K lower for the short wave and long wave bands respectively compared to the previous cases. This reduction in apparent temperature is

mixing with the cooler ambient air by increasing the plume perimeter, thus, increasing the region influenced by turbulent mixing.
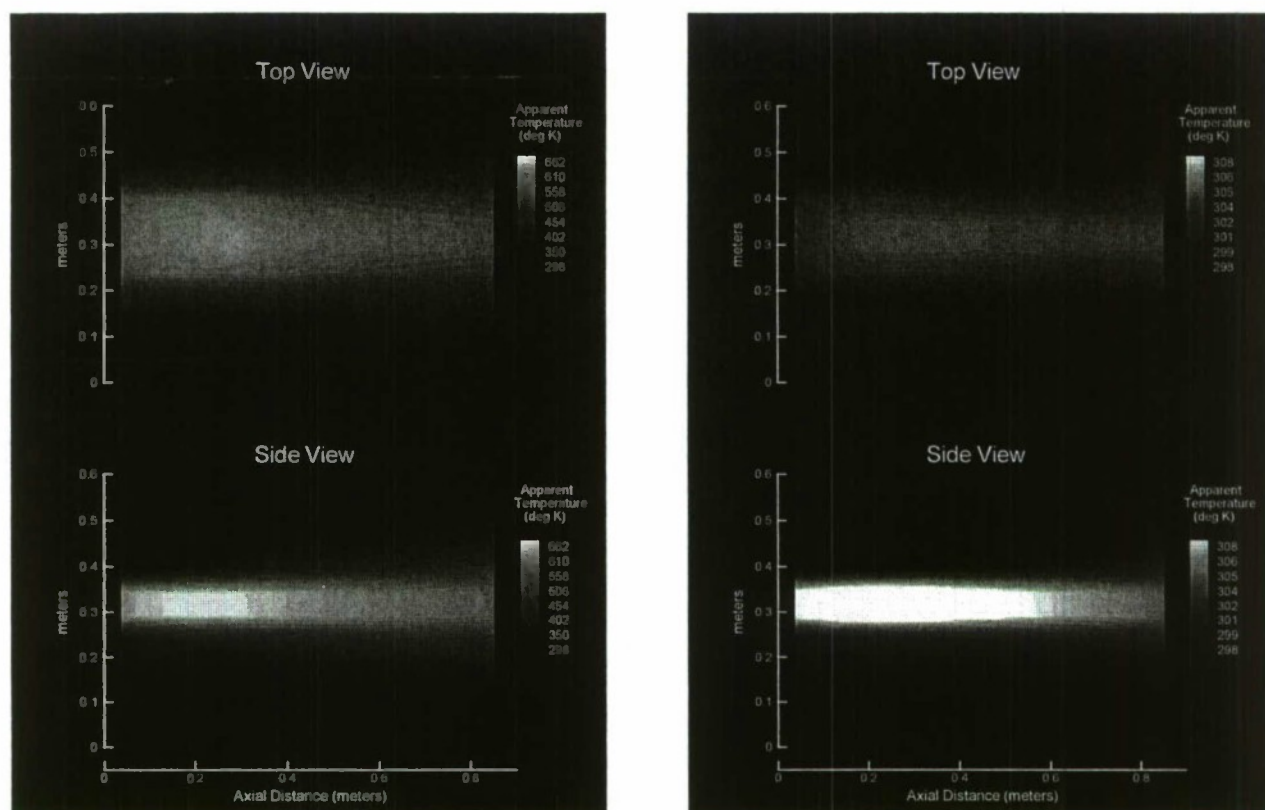
Figure 6: Contours of apparent temperature of a rectangular cross section exhaust with a 2.0 m/sec cross wind. Spectral ranges 3-5 micron (left) and 8-12 micron (right).

CONCLUSION

This paper reviewed recent work used to estimate the thermal signature of an exhaust plume. The method incorporated a CFD model to generate the temperature and species concentration throughout the plume structure. Custom written software discretized the CFD derived solution into small control volumes. MODTRAN was applied to each volume, thus, providing estimates of the absorption and emissive power within each control volume. The software integrated the effects of the local species along a user-defined line of sight. For a given line of sight, the resulting two-dimensional image consisted of contours of apparent temperature.

REFERENCES

1. Fair M.L., 1996, "NATO Infrared Air Target Signature Prediction Model (NIRATAM)", in 2nd NATO/IRIS Joint Symposium, London.

2. Morin J., F. Reid, and D. Vaitekunas, "SHIPIR: a model for simulating infrared images of ships at sea", SPIE 2223, 1994.

3. Berk, A., L.S. Bernstein, and D.C. Robertson, "MODTRAN: A Moderate resolution Model for LOWTRAN 7", AFGL-TR-89-0122, 1989.

SMOKE AND OBSCURANT MODELING IN SUPPORT OF SIMULATION BASED ACQUISITION AND TRAINING

David J. Johnston
Bruce W. Fischer
OptiMetrics, Inc.
Forest Hill, Maryland 21050
Telephone: (410) 893-9714
Telefax: (410) 893-9717
E-Mail: johnston@omi.com

William G. Rouse
U.S. Army Edgewood Chemical Biological Center
AMSSB-RRT-PA
Aberdeen Proving Ground, Maryland 21010
Telephone: (410) 436-1848
Telefax: (410) 436-2998
E-Mail: wgrouse@sbccom.apgea.army.mil

ABSTRACT

The U.S. Army Edgewood Chemical Biological Center (ECBC) and its parent command are committed to the implementation of simulation-based acquisition and training techniques. Their value has been demonstrated repeatedly by many organizations and in countless situations representing all aspects of the product life cycle, including combat development, material development, manufacturing, training, and employment. Models and simulations have become practical tools that routinely replace laboratory experiments, field activities, and training exercises while reducing cost and improving safety. They could be particularly effective for smoke and obscurant scenarios, which carry significant environmental clearance costs and additional safety considerations. Numerous models have been developed to simulate smoke and obscurant systems, but most of these have focused on detailed scientific and engineering issues, such as cloud physics and electro-optical performance. Although several attempts have been made to integrate these models into constructive and distributed interactive simulations, no smoke and obscurant model has emerged

that supports operational level activities. This is a significant shortcoming that hampers both training and acquisition. This paper reviews the state of smoke and obscurant models and illustrates how existing simulations have been used to support acquisition and training. It also identifies requirements for operational-level tools and establish a framework for their development.

1. MODELING AND SIMULATION TRENDS

Modeling and simulation (M&S) has become an essential technology that is influencing a growing number of mission-critical activities. Organizations within DoD are increasingly using it for operational planning; research, development, and acquisition; test and evaluation; training and mission rehearsal; and doctrine development. This growth is being fueled by budgetary constraints, environmental concerns, and technological advances, which reduce the cost of complex computer systems while increasing their capabilities.

Historically, most M&S activities have been conducted in vertical stovepipes within application domains (figure 1)

Analytical Simulations	Operational Level Training Simulations	Tactical Level Training Simulations	Training Ranges	Real Weapon Systems and C4I	Test and Evaluation Ranges	Engineering Level (R&D, T&E) Simulations	Manufacturing Simulations	Other Simulations
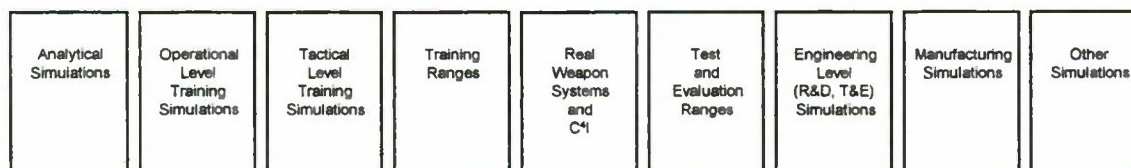
Figure 1. DoD modeling and simulation domains.

and there has been little horizontal integration. The constructive wargames and virtual simulations used for training have seldom been linked, for example, to engineering-level simulations or the physical models they use.

The computational burden imposed by high fidelity models has often limited their utility beyond the application domain for which they were developed. This is certainly true for smoke and obscurant models, which have little presence in the preeminent simulation environments. Until recently, horizontal integration was infeasible due to limitations imposed by hardware, software, and design methodologies. This situation is changing, however, as technological developments increase computer system performance, simplify connectivity, reduce cost, improve usability, enhance the development process, and promote software reuse. These factors are reshaping the nature of DoD models and driving a paradigm shift toward a more flexible and efficient simulation environment (figure 2).

The Defense Modeling and Simulation Office (DMSO) has capitalized on these trends to produce a High Level Architecture (HLA)[1] that unifies all M&S domains under a single technical framework. Future DoD M&S systems will be required to comply with the HLA, which specifies a design philosophy, imposes documentation requirements, and provides a common cross-platform run time infrastructure (RTI). The HLA will facilitate linkages between models and promote software reuse across all M&S domains. We are entering a period where software, that was developed for one specific purpose, can be used without modification in many different applications. This will enable engineering-level, constructive wargames, and virtual simulations to use physical models to increase their fidelity with real world phenomena.

2. SMOKE AND OBSCURANT MODELS

Several physical models have been developed to simulate the production, transport, and diffusion of battlefield obscurants and assess their effect on tactical sensors. The U.S. Army Research Laboratory (ARL) has been a major contributor and it maintains two smoke and obscurant models in its Electro-Optical Systems Atmospheric Effects Library[2,3]. GRNADE[4] simulates multiple-round salvos of tube-launched grenades (L8A1 and M76) and is used for self-screening analysis. The Combined Obscuration Model for Battlefield-Induced Contaminants (COMBIC)[5,6] is more comprehensive and can simulate: high explosive and vehicular dust; phosphorus and hexachloroethane munitions; diesel fuel, oil, and rubber fires; generator-disseminated oils; other screening aerosols, and user-defined sources.

COMBIC has been used in numerous, diverse applications and is arguably the dominant model in this field. It operates on level terrain and only considers a horizontally homogeneous wind (figure 3). These limitations led ARL to develop a derivative model, the Simulation of Aerosol Behavior in Realistic Environments (SABRE)[7], that can use a terrain-dependent wind field. SABRE was an EOSAEL module for some time, but is no longer supported by ARL and has been withdrawn from the library, leaving it without a smoke and obscurant model that handles non-uniform wind and terrain.

The Joint Project Office for Smoke, Obscurants, and Special Technologies Counter-measures (JPO-SOSTC), Naval Surface Warfare Center (NSWC), ECBC, and OptiMetrics, Inc. (OMI) have enhanced SABRE to create the Transport, Diffusion, and Radiance (TDR) model[8]. It has been integrated into several applications and can operate in a stand-alone mode on numerous computing platforms.

Stand-alone simulations	Networked simulations
Tailored-made code	Reusable code
Procedure-oriented code	Object-oriented code
Module interaction via procedure calls	Module interaction via messages
Tightly coupled architecture Predetermined scale Not easily reconfigured	Loosely coupled architecture Undetermined scale Rapidly reconfigurable (plug and play)

Figure 2. Modeling and simulation trends.
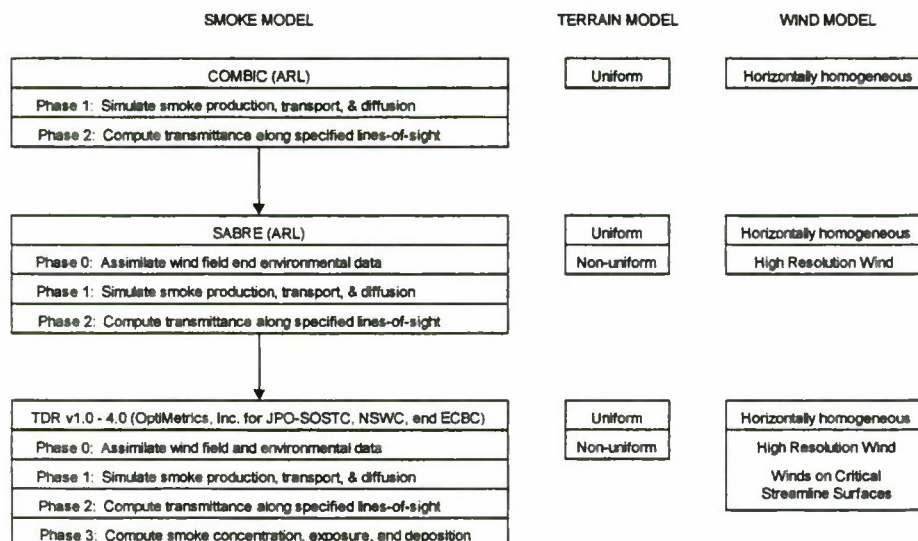
Figure 3. Lineage and primary characteristics of COMBIC and its derivatives.

COMBIC and TDR were originally developed in the mid 1980s to the early 1990s using prevailing techniques. They are large, monolithic programs that are written in FORTRAN using extremely unstructured code. Their software components make extensive use of global variables and are, therefore, highly interdependent. In short, they are not well suited for use in a simulation environment that is increasingly distributed and object-oriented.

Both models describe clouds as a collection of one to five subclouds, which can be either Gaussian puffs or plumes. When a source is activated, subcloud states are computed at discrete downwind distances and the results are saved in a history file. All predictions are made using the atmospheric conditions that existed at the time of source activation. Although this approach is computationally efficient, it is not responsive to changes in atmospheric conditions that might occur during the cloud's lifetime. It also requires a large amount of data to be maintained (and possibly transferred) for subsequent calculations.

Because they were developed several years ago, COMBIC and TDR only simulate sources and obscurants that were available (primarily to U.S. forces) at that time. The models have not been updated significantly to include equipment, munitions, or materials that have been fielded in recent years or are currently under development by the U.S., our allies, and potential adversaries. Both models enable the user to specify source and obscurant characteristics through data inputs, but this requires an intimate knowledge of the material and model properties.

Neither of these programs model vehicles or vehicular components. They have no awareness of specific vehicle types nor the location and orientation of smoke generation equipment on those vehicles. Consequently, COMBIC and TDR could not be used by themselves to examine operational usage where component placement is an issue. This limitation is aggravated by their inability to accept unscripted inputs. The models cannot be used without augmentation to respond to ad hoc smoke events that might be generated randomly, by a constructive wargame under player control, or by networked simulators in a virtual exercise. In addition some deficiencies have been noted in their predictive algorithms. Most notably, COMBIC does not accurately model evaporative losses from disseminated oils as a function of temperature[9]. This affects the predicted quantity of suspended liquid that is actually available for producing screening effects, a critical factor in some applications where oil smokes are employed.

3. SMOKE SYSTEM PERFORMANCE MODEL

The Smoke System Performance Model (SSPM)[10] was developed by ECBC and OMI to eliminate many of the limitations noted above. It is a collection of C++ classes that model the essential elements of smoke and obscurant systems. The classes can be integrated with engineering-level models, constructive wargames, and virtual simulations to enhance their ability to simulate battlefield obscuration.

SSPM models relevant items, such as vehicles, components, clouds, obscurant materials, and vehicular grenades. Each class encapsulates the essential technical characteristics of the item it represents, as they relate to smoke and obscurant production, and provides default functionality. The default behaviors vary in sophistication

and can be overridden if enhanced functionality is required. The notional default behavior of SSPM clouds can be enhanced by using a smoke and obscurant model, such as COMBIC or TDR, to simulate cloud production, transport, and diffusion. When this is done, SSPM acts as a preprocessor by simulating operations at a higher level and directing the smoke and obscurant models to place clouds with specified characteristics at designated times and places.

SSPM is limited only by the scope of the systems it currently models. The latest version can simulate thirteen vehicles, seven vehicle-launched grenades, vehicle engine exhaust smoke systems, and two smoke generators. It does not yet model smoke pots or artillery, mortar, rocket, and aircraft-delivered obscurants. Only one of the vehicles and one of the grenades are foreign systems.

4. ENGINEERING LEVEL MODELS

SSPM has been linked to COMBIC in the Cloud Density Visualization Utility (CDVis)[10], an engineering-level model that presents a graphical representation of simulated clouds (figure 4). CDVis uses SSPM to execute complex obscuration scenarios, COMBIC Phase I to predict cloud histories, and COMBIC Phase II to compute concentration path lengths along specified lines-of-sight. The concentration path lengths or their corresponding transmittance values are then presented as false color

images, enabling the user to perceive the clouds in three dimensions as a function of time.

SSPM has also been linked to COMBIC in a battle management system (BMS) for chemical staff officers that will be used to evaluate smoke and obscurant plans (figure 5). The BMS is similar to CDVis in its use of SSPM and COMBIC, but it superimposes a birds eye view of the simulated clouds on a tactical land map. The BMS also predicts sensor effectiveness on a horizontal plane from a given location at a designated time and distance above ground level. Effectiveness is presented as a radar plot that uses green, amber, and red to depict regions of increasing cloud density in accordance with specified transmission thresholds.

This program demonstrates the power of object-oriented methodologies, which foster software reuse by encapsulating functionality and enabling components to be assembled into new and different applications. Most of the smoke-related code in BMS is identical to that used in CDVis. The only difference is some minor additions that were applied to support its unique display requirements. Also, the BMS graphical user interface is written in Java while SSPM and its CDVis extensions are written in C++ and the extended cloud class places an object wrapper around COMBIC's FORTRAN code. The object-oriented technology enables these disparate components to be drawn together with relative ease to create a new and useful application.
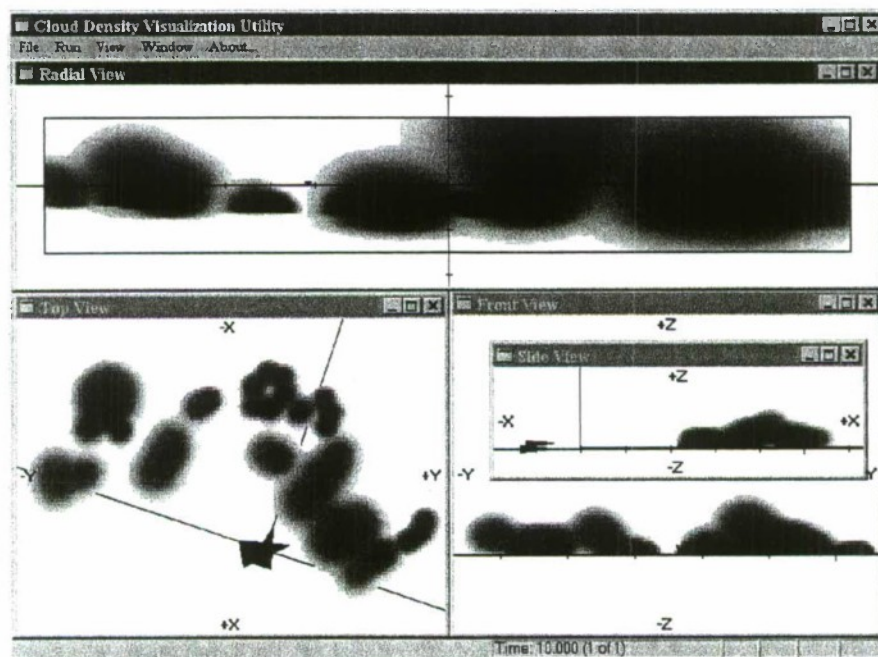
Figure 4. CDVis displaying cloud density images for the front, side, top, and radial views of a simulated obscuration event.
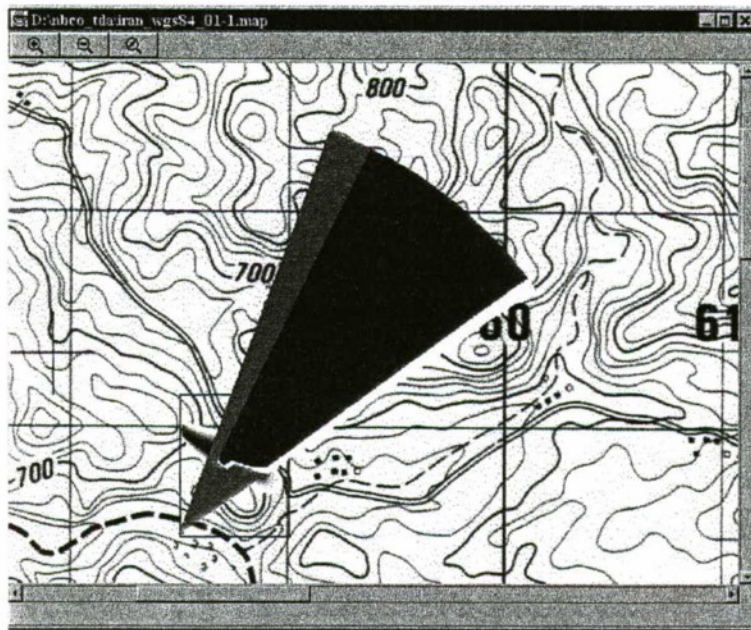
Figure 5. Chemical Staff Officer's BMS displaying a sensor effectiveness plot
superimposed on a cloud density image of a simulated obscuration event.

5. CONSTRUCTIVE WARGAMES

JPO-SOSTC, NSWC, the USMC Systems Command, and OMI have integrated TDR into a many-on-many Sensor/Obscurant Engagement Simulation (SOES). It has been used by combat and material developers to assess tactical concepts for smoke employment in littoral operations (figure 6). SOES uses physical models to simulate sensor performance and TDR Phase I to produce terrain-sensitive smoke clouds. Intervisibility issues are addressed using digital terrain elevation data and TDR Phase II.

SOES uses TDR as an independent executable program without augmentation, so it is limited to the smoke sources that TDR inherently supports. These sources are positioned and activated in accordance with an scripted scenario.

The US Army Training and Doctrine Command has integrated COMBIC into CASTFOREM, a stochastic force-on-force model that is used to assess combat system performance[11]. It is the Army's primary tool for conducting formal Analysis of Alternatives (AoA). CASTFOREM uses COMBIC to predict transmissivity through obscurant clouds, which have been produced by simulated combat events. When activated, COMBIC is employed throughout the gaming exercise, but the smoke effectiveness assessments are limited to discrete engagement segments and only affect certain operations, such as laser range finding and missile flyout. During laser operations, CASTFOREM does consider the affect of obscurants on a laser beam and will attenuate the returning signal appropriately. During guided missile flyouts, CASTFOREM periodically uses COMBIC to determine if obscurants have cause the missile to break lock.

Obscurant usage places a large computational burden on CASTFOREM and can significantly increase run times. Because COMBIC is used in its native form, CASTFOREM can only use standard sources. And, because there is no direct linkage to operational entities, such as vehicles, artillery units, and aircraft, smoke events must be manually inserted by exercise controllers.

6. VIRTUAL SIMULATIONS

In recent years, DoD has invested heavily in distributed interactive simulation (DIS) technologies that enable manned and unmanned simulators to interact on a common virtual battlefield. DIS simulations have been used extensively for small unit training and are increasingly used in other applications, including simulation based acquisition. Interactions are facilitated through the exchange of messages using well defined protocols. The DIS standard provides for the transmission of some smoke information, enabling one battlefield entity to produce a smoke event and report its state to all other networked simulations.

The Modular Semi-Automated Forces (ModSAF) program is used extensively in DIS applications to populate the virtual world with battlefield entities, such as
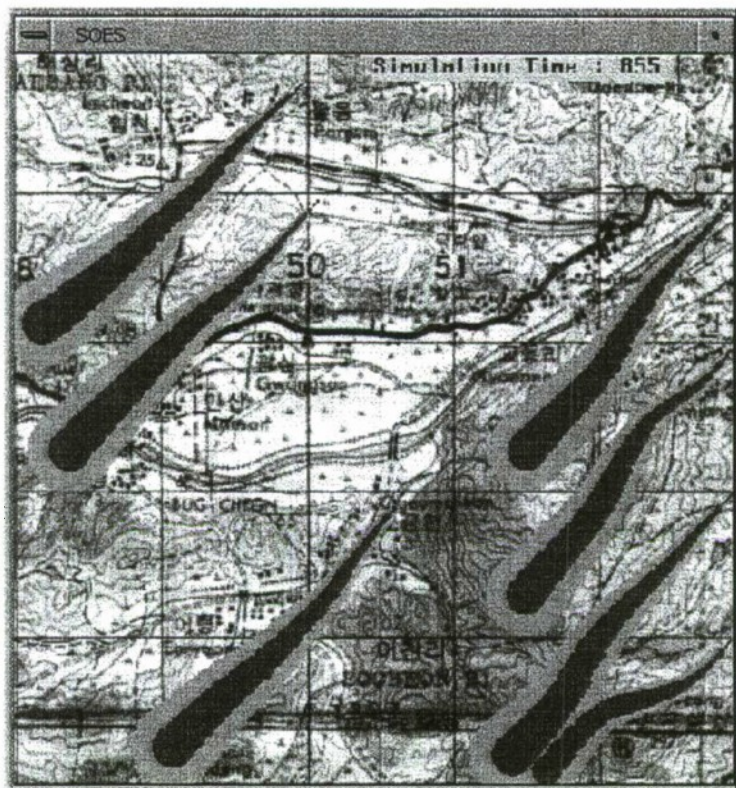
Figure 6. A SOES display depicting terrain sensitive TDR-generated clouds.

aircraft and armored vehicles. These entities behave in an intelligent manner and are indistinguishable from their manned counterparts.

The use of smoke and obscurants by ModSAF entities is notional. Certain vehicle types have a limited ability to launch salvos of self-protective grenades and artillery/mortar projectiles. No other smoke sources are supported. ModSAF uses pre-computed COMBIC history files to instantiate obscurant clouds, but these clouds are limited to one subcloud each, which severely restricts their fidelity. ModSAF does consider obscurant effects on entity engagements, but few other DIS applications do. The scene generators on most manned DIS simulators cannot render smoke clouds (with the exception of some trailing effects attached to some entities) and they do not affect crew vision whatsoever.

The Close Combat Tactical Trainer (CCTT) is a family of networked simulations that is used to train armor, cavalry, and mechanized infantry platoons in the performance of collective tasks. CCTT includes manned simulators for numerous combat vehicles and a semi-automated forces program (similar to ModSAF) that can control a wide variety of friendly or opposing units. CCTT uses a variant of the DIS standard protocols to establish and maintain the synthetic environment.

CCTT has an extremely limited capability for simulating smoke events and only supports three obscurant types: hydrochloric acid, red phosphorus, and white phosphorus. Manned CCTT simulators do render smoke clouds using an animation technique that can vary transmittance in accordance with obscurant characteristics. This does affect crew visibility which can influence combat operations. Other obscurant effects are not supported.

7. FUTURE REQUIREMENTS

Smoke and obscurants do affect battlefield sensors and those effects can influence combat operations. It is important for obscurant systems to be reasonably represented in models and simulations so that their influence can be properly assessed. The capabilities and limitations of U.S., allied, and opposing forces must all be considered. Smoke and obscurant modeling should not be done for its own sake, but it must be done to insure that soldiers are properly trained and equipped to operate on the dirty battlefield. It is incumbent upon the smoke and obscurant community to make sure that occurs.

The survey presented above describes the current state of smoke and obscurant modeling. It is by no means comprehensive, but does highlight strengths and weaknesses in several application domains. The survey illustrates that smoke and obscurant modelers have done

227

have developed excellent physical models that reasonably simulate the growth, transport, and diffusion of obscurant

requirements and computational constraints. And, the physical models have been used in numerous applications

survey also illustrates, however, that smoke and obscurant modelers have failed to accomplish what heretofore was infeasible to do. Despite all good intentions and a lot of intense effort, they have not managed to insert a significant amount of smoke and obscurant play into the tactical simulations that are routinely used for training and simulation-based acquisition. Given technological limitations, it was just too difficult to achieve. That situation is changing.

Recent advancements in hardware and software technologies are enabling simulations to model physical phenomena with increasing fidelity. The emergence of new design methodologies is facilitating the development of true software components that can readily be used in diverse computing environments. This is an excellent time to begin the development of smoke and obscurant components for the simulation community at large.

These components must be comprehensive, flexible, authoritative, efficient, self-contained, and HLA compliant, as described below:

Comprehensive. Collectively, the simulation components should model all smoke and obscurant systems that U.S. forces are likely to use or encounter on any future battlefield.

Flexible. For maximum applicability across all M&S domains, each component (or variation thereof) must be able to operate at several resolutions. High fidelity simulations will need high fidelity smoke representations while low fidelity simulations will need the opposite. Also, cloud behavior must be 4-dimensional (i.e., sensitive to spatial and temporal variations in atmospheric conditions, if they exist).

Authoritative. Each component must simulate the item it represents with sufficient fidelity to satisfy the smoke and obscurant community at all possible resolutions.

Efficient. Smoke and obscurant modeling is computationally intensive and that burden has limited its utility in many applications. The simulation components must employ more efficient algorithms than those used today, particularly for modeling cloud behavior and obscurant effects. The use of neural nets and related technologies should be

performance without degrading fidelity. Also, data exchange requirements among networked simulators

Self-contained. Developers will not use the applications. Consequently, they must be designed using accepted object-oriented techniques, which require them to encapsulate all characteristics and behaviors and expose a well-defined interface.

HLA compliant. The simulation components must be developed in accordance with the HLA specification. They must be fully documented with federation object models (FOM) and/or simulation object models (SOM), as appropriate. The components must employ the simulation support services provided by the RTI.

8. REFERENCES

1. Defense Modeling and Simulation Office, 1995: HLA Management Plan: High Level Architecture for Modeling and Simulation, Version 1.6. Defense Modeling and Simulation Office, Alexandria, Virginia.

2. Shirkey, R.C., L.D. Duncan, and F.E. Niles, 1987: EOSAEL 87, Volume 1, Executive Summary. U.S. Army Laboratory Command, Atmospheric Sciences Laboratory, TR-0221-1, White Sands Missile Range, New Mexico.

3. Mason, J.B. and R.G. Steinhoff, 1987: EOSAEL 87, Volume 2, Executive User's Guide. U.S. Army Laboratory Command, Atmospheric Sciences Laboratory, TR-0221-2, White Sands Missile Range, New Mexico.

4. Davis, R.E. and R.A. Sutherland, 1987: EOSAEL 87, Volume 14, Self-Screening Applications Module GRNADE. U.S. Army Laboratory Command, Atmospheric Sciences Laboratory, TR-0221-14, White Sands Missile Range, New Mexico.

5. Hoock, D.W., R.A. Sutherland, and D. Clayton, 1987: EOSAEL 87, Volume 11, Combined Obscuration Model for Battlefield-Induced Contaminants (COMBIC). U.S. Army Laboratory Command, Atmospheric Sciences Laboratory, TR-0221-11, White Sands Missile Range, New Mexico.

6. Ayres, S.D. and S. DeSutter, 1993: Combined Obscuration Model for Battlefield Induced Contaminants (COMBIC92) User's Guide (Draft). U.S. Army Research Laboratory, Battlefield

Environment Directorate, White Sands Missile Range, New Mexico.

7. Maynard, H.W., D.W. Hoock, F.W. Hansen, T. Henmi, J.P. Cruncleton, and W.D. Ohmstede, 1987: EOSAEL 87, Volume 12, Simulation of Aerosol Behavior in Realistic Environments (SABRE). U.S. Army Laboratory Command, Atmospheric Sciences, TR-0221-12, White Sands Missile Range, New Mexico.

8. Pennsyle, R.O. and F.J. Wysocki, 1998: Transport, Diffusion, and Radiance (TDR) Version 4.0 Users Manual. OptiMetrics, Inc. for the U.S. Army Edgewood Research Development and Engineering Center, Naval Surface Warfare Center, and Joint Program Office for Special Technology Countermeasures, OMI-639, Forest Hill, Maryland.

9. Pennsyle, R.O., 1996: Droplet Evaporation in Oil Smokes. Proceeding of the Smoke/Obscurant Symposium XIX, Laurel, Maryland.

10. Johnston, D.J. and W.G. Rouse, 1999: The Smoke System Performance Model and Cloud Density Visualization Utility Version 2.0. OptiMetrics Inc. for the U.S. Army Edgewood Chemical Biolgical Center and the U.S. Army Program Manager for Smoke and Obscurants, OMI-653, Forest Hill, Maryland.

11. Horton, W.D. and J.A. Seton, 1999. Hit Avoidance Analysis, Test, and Combat Effectiveness Assessment, Draft Final Report. OptiMetrics Inc. for the U.S. Army Research Laboratory and the Office of the Test Director, Precision Guided Weapons Countermeasures Test and Evaluation Directorate, OMI-640, Las Cruces, New Mexico.

A Robust Data Set for Visual Detection Model Calibration and Validation
for the Proceedings of the
Ground Target Modeling and Validation Conference

R. D. Ellis
Wayne State University
87 E. Ferry St
Detroit, Ml 48202 USA
E-mail: rdellis@wayne.edu

G. Witus
Turing Associates, Inc.
1392 Honey Run Drive
Ann Arbor, MI 48103 USA
E-mail: witusg@umich.edu

ABSTRACT

This paper describes an effort underway to create a large image set and corresponding target acquisition performance data for use in the calibration and validation of computational models of human visual processing in target detection tasks. A set of 1154 images was created by digitally altering a base set of 44 images of nine types of vehicles in varied aspects and locations. The objective of the digital modification was to create systematic variation in luminance, color, contrast, size scale, and cue features. A target detection experiment with 21 human observers was used to collect human response data. Assessment of the data indicated that the digital manipulations introduced a greater amount of overall reliable variance in observer response.

INTRODUCTION

The past decade has seen the computational approach to modeling the human observer response advance and enjoy continued popularity. Theory has developed rapidly in the area of human visual processing, and this has led to the development of more and more sophisticated algorithms and corresponding computational implementations. A limiting factor in the wider acceptance of such models is the difficulty involved in empirical validation and calibration using human observer data.

Currently, the validation and calibration (V/C) efforts of various models are relatively independent. A more integrated approach to V/C, using a common set of images and corresponding human response data would allow assessment and comparison of results across models, and would also reduce duplication of effort involved in *ad hoc* collection of human response data.

In June of 1999 the NATO SCl-12 Working Group on Camouflage, Concealment and Deception Evaluation Methodologies held a workshop on Search and Target Acquisition (STA) Modeling. Participants in the workshop were provided with a copy of the TNO Search_2 image set, and the corresponding search performance data. The image set consisted of 44 images of 9 different types of military vehicles, taken at Hunter-Ligget, at ranges from 500 to 5,000 meters, at varying aspect angles and slide film camera using a 10x lens, then digitized at 6K by 4K resolution. The search performance data (response time and probability of detection) was collected in laboratory testing with 62 subjects using the projected slides. This data set provided a common basis to compare alternative search and target acquisition modeling and analysis methods.

During the workshop, the researchers noted several characteristics of the data set that limited its usefulness for STA model calibration and validation:

1. The vehicles did not contain significant variation in color, reflectivity or camouflage. The camouflage patterns painted on the vehicles were not discernible or effective at range.

2. The images did not contain significant variation in overall scene luminance and contrast trasmittance.

3. Most of the vehicles were easy to detect. Only 3 of the 44 targets had probability of detection (Pd) below 0.5, and only 4 of the 44 had mean search time above 20 seconds. 75% of the targets had Pd above 0.8 and search time below 10 seconds.

In short, the Search_2 stimuli did not contain sufficient variation, in either appearance or human response, to provide a robust calibration or validation of STA models.

There are only a few widely available unclassified image sets that have accompanying human performance data, such as Search_2, and those that are available are not adequate for a robust calibration and validation. The use of a large data and image set, with widely varying image properties and human responses, would facilitate better model validation and calibration. Use of a robust common data/image set would facilitate model comparison, and would also reduce the burden of the V/C process on developers. This paper describes the development of a large data and image set with widely varying image properties and reliable human performance data.

STIMULUS GENERATION

A set of 1154 images was derived by digitally altering the 44 Search_2 images. The digital alteration followed 2 main principles: First, images were altered to produce structured variations in scene and target appearance. Second, alterations were performed to introduce greater range and variation in the human response variables, Pd and Td.

These images were then cropped to 1080 x 720. The images were cropped so that no two images presented the identical scene, and so that the targets were equally distributed in the four quadrants of the image. Half of the images were flipped horizontally to further disrupt possible scene memory effects. The images were stored in high-resolution JPEG format.

Image background characteristics were modified in the expanded image set with regard to overall color/luminance and image resolution. The color/luminance variations were performed as follows: 25% of the images were left at baseline (unaltered color/luminance appearance), 25% were darkened (50% brightness reduction), 25% were given artificial "haze" (50% contrast reduction and 50% brightness increase), and 25% were transformed into gray-scale (the color was removed). Resolution (i.e., a proxy for range/size scale) was manipulated as follows: 50% of the images were sampled at 1/2 the resolution of the Search_2 images and 50% of the images were sampled at 1/3 the resolution of the Search_2 images.

Target appearance characteristics were modified in the expanded image set as well: in 27% of the images the baseline target appearance was unchanged, in 27% the targets were modified to reduce the target-to-background contrast by one-third, in 19% the targets appearance was modified to suppress specific cue features (e.g., by darkening glint facets, overlaying camouflage nets, etc.), and in 27% of the images the target was digitally replaced with an empty background to serve as "no target" controls.

METHODS AND PROCEDURES

Stimulus Presentation and Apparatus.

The observer task in this experiment was detection given a cue to the target location. The observer was cued to the target location by a thin red circle approximated three degrees in diameter centered on the target. In the "no target" control images the circle was centered on a potential target location. (A subsequent narrow field of view search experiment is planned, using the same images but without the target location cue).

The observers were instructed to respond to the question ("Is there a vehicle in the image?") using a four-choice menu:

  1 – a vehicle was DEFINITELY present;

  2 – a vehicle was PROBABLY present;

  3 – a vehicle MAY BE present; and

  4 – a vehicle was DEFINITELY NOT present.

Response time and perceived target location (i.e., mouse click) were also recorded. The stimuli remained on the screen until the observer entered a response. There was no artificial response time window.

Stimuli were presented on a Gateway2000 Pentium II with a 17" EV700 display, driven by an ATI Rage graphics card set to 24-bit color, 1280 X 1024 resolution. A custom stimulus presentation and data acquisition program was written in Microsoft Visual Basic 6.0, with a third-party add-on DLL for high-resolution, high-color JPEG presentation.

Procedure.

Twenty-one university students participated in this phase of experimentation. They were all screened for 20/20 near acuity and color vision problems. Observers were seated 60" from display and presented with the imagery. At this distance, each pixel subtended 0.01 degrees (well below visual spatial resolution).

Observers reported on each of the 1154 images twice, once according to the manipulations described above and again at 50% scale (which was handled on-the-fly by the presentation software to further increase the variability in appearance). At full scale, the images subtended 10.5 degrees by 7 degrees.

Observers responded to each presentation with a mouse, by clicking an item on 4-point forced-choice rating scale (the item labels ranged from NO vehicle present to DEFINITELY a vehicle present).

Trial order was determined by assigning stimuli to 4 experimental blocks via a Latin square. Blocks were replicated across presentation scale (8 total blocks). Block order and stimulus order within block were randomized.

Data Collection and Treatment.

The independent variables recorded included the original image number, target treatment (baseline, low contrast, special variation, and no target), scene treatment (baseline, darkened, haze, and grey-scale), image sampling resolution (1/2 and 1/3 original resolution), and presentation scale (full and half). Dependent variables included the participants' response on the 4-point rating scale, response time, and mouse position (XY coordinates in pixels).

The rating scale data were compared against ground truth data to generate three hit rate (HR) and false alarm rate (FAR) pairs, which were then used in turn as three different points on a receiver-operating characteristic (ROC) curve for each of the 1154 images. The three ROC points corresponded to different response biases. A HR/FAR pair corresponding to a conservative response bias was computed by counting "Definite" responses as a target present response. A moderate bias was computed by aggregating the "Probably" and "Definitely" responses, and a liberal bias was computed by counting "Maybe", "Probably", and "Definitely" against "Nothing."

DATA ASSESSMENT

The primary objective of this project was to create a large, varying image set with reliable and detailed human response data corresponding to each image. Assessment of the data on a descriptive level indicates that the objective was accomplished.

Aggregating all manipulations by image number, the new image set had wide overall variation in observer ratings (ranging from 0.7 to 2.9 on a 0-3 scale), and did not have ceiling or floor effects present in the TNO Search_2 data.

The scene treatments chosen, on average, created a pattern of results consistent with expectations. Compared to baseline

scene characteristics, the effect of haze, darkening and eliminating color information all served to reduce observer ratings (they were less sure that targets were indeed present in target-present images). Target properties were consistent as well; reducing contrast, eliminating glints, etc., resulted in lower ratings.

For example, Figure 1 depicts the ROC curve for two different target treatments for Image 38. The baseline target is shown in Figure 2, while Figure 3 depicts the same target with glint removed (Note: these figures were further manipulated for presentation in this paper, and did not appear to observers as they do in the figures). Using area under the ROC curve as a measure of observer sensitivity, it is easy to see that removing the glint (while that is a relatively small effect on overall target contrast) has a great effect on sensitivity. Note also that reducing the overall target contrast by 50% (not depicted) reduced target visibility, but not nearly as much as removing the glint did.
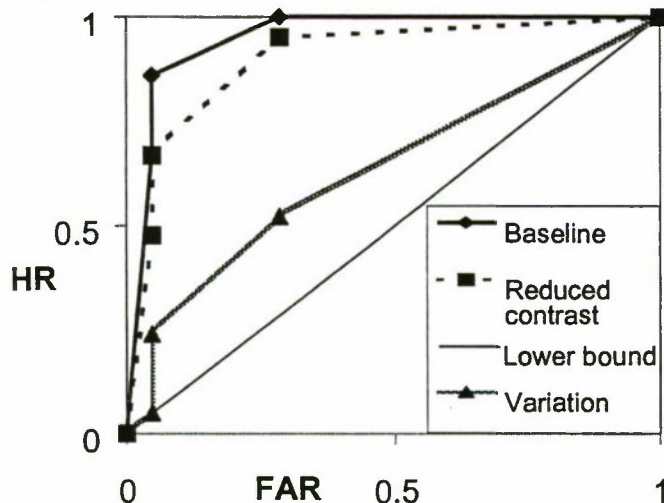
Figure 1. ROC curve for Image 3. Each point represents 21 observer responses.

Figure 2. Image 38 with baseline target

Figure 3. Image 38 with target glint suppressed

DISCUSSION & CONCLUSIONS

Overall, the results from the experiment were very encouraging and consistent with expectations. The image set contained a wide variation in scene and target characteristics. Responses across observers were tightly grouped for each image, indicating low observer variance. The observer responses contained wide variation from image to image. These variations have not been analyzed in detail, but are generally consistent with the type of image manipulation (e.g., darkening, reducing contrast transmittance, reducing target contrast and removing color all reduced target detectability), and with previous test results using only the baseline imagery.

Taken together, the image set and the perception test results constitute a large and robust data set for calibrating target acquisition models. They contain a wide variation in observer response, and systematic variation in image and target appearance.

A search experiment on the same image set is underway. The phase II study will use the same independent and dependent variables, but stimuli will be presented without an annulus or other cues as to target location. This study will collect data from 100 observers, screened for good vision and trained with regard to vehicle properties and response criteria.

It is expected that the entire image set and accompanying search and detection data will be available for model validation in the Spring-Summer 2000 time-frame.

ACKNOWLEDGEMENTS

REFERENCES

Toet, A. et al., A High-Resolution Image Data Set for Testing Search And Detection Models, *TNO-Report TM-98-A020*, Soesterberg, The Netherlands: TNO Human Factors Research Institute, 1998.

232

Applying the Law of Comparative Judgement to Target Signature Evaluation

James R. McManamey
U.S. Army Communications and Electronics Command
Night Vision and Electronic Sensors Directorate
10221 Burbeck Road, Suite 430, Building 305
Fort Belvoir, Virginia 22060-5806
E-mail: jmcmanam@nvl.army.mil

ABSTRACT

The Law of Comparative Judgement (LCJ) is a psychophysical tool that can be used to scale complex phenomena that lack easily identified physical parameters. Target signatures represent such phenomena. In a demonstration exercise, a "search difficulty" value was found using the LCJ. These LCJ scale values were compared to search times and probabilities of detection from a search experiment run in the Netherlands. The scale values were not linearly related to search time and probability of detection, but correlated very well with the logarithm of mean search time (r = 0.936) and the cube of the number of correct responses (r = 0.954). A chi-squared goodness-of-fit test gave 94.6% confidence in the fit of the LCJ scale to the experimental data. While the LCJ results in a scale with no natural zero point and arbitrary units, this tool can be used to construct a standard scale. This paper illustrates how a standard clutter scale might be constructed using the LCJ. The LCJ could be a valuable tool in target signature evaluation either when used in conjunction with scaling equations that permit conversion to familiar quantities such as mean search time and probability of detection, by providing relative "search difficulty" values, or by making possible a psychophysically meaningful clutter scale.

INTRODUCTION

Today, there are many quantities that engineers and scientists want to measure in perceptually meaningful ways. For example, designers of military man-in-the-loop search and target acquisition systems, as well as engineers working on military signature suppression systems, want measures of effectiveness that are psychophysically meaningful, repeatable, and correlate well with field performance. Such measures of effectiveness have frequently been surprisingly elusive. Target detectability and signature levels may seem like concrete, physically measurable quantities, but in truth they have much in common with such abstract concepts as beauty.

Figure 1 shows a near-infrared scene. The upper image shows a tank profile that has been inserted into the scene. In the lower image, the tank is not visible at all. It is "perfectly camouflaged." However, most signature evaluation models and virtually all of the most widely used sensor models would say that the two tanks have exactly the same signature. This is because the only difference between these two target signatures is that the image pixels have been moved around. Averaged over the target, the histogram, contrast, variance, third-, fourth-, and fifth-moments are all the same. Only measures of effectiveness that can distinguish between the relatively large "blobs" in the lower image and the "salt-and-pepper" noise in the upper image can distinguish between the two tanks. Only a model that can determine that the tank in the lower image has the same "texture" as the background and that the edges are perfectly "blended" with the background while, at the same time, determining that these things are not true of the tank in the upper image, can accurately predict that a person will detect the target in the top picture and fail to detect the target in the lower one.

Investigators around the world are trying to develop models that can make such distinctions. Many of these models, a type called "computational vision models," attempt to mimic various processes that are believed to take place in the human eye-brain system. This has been a daunting task, and none of
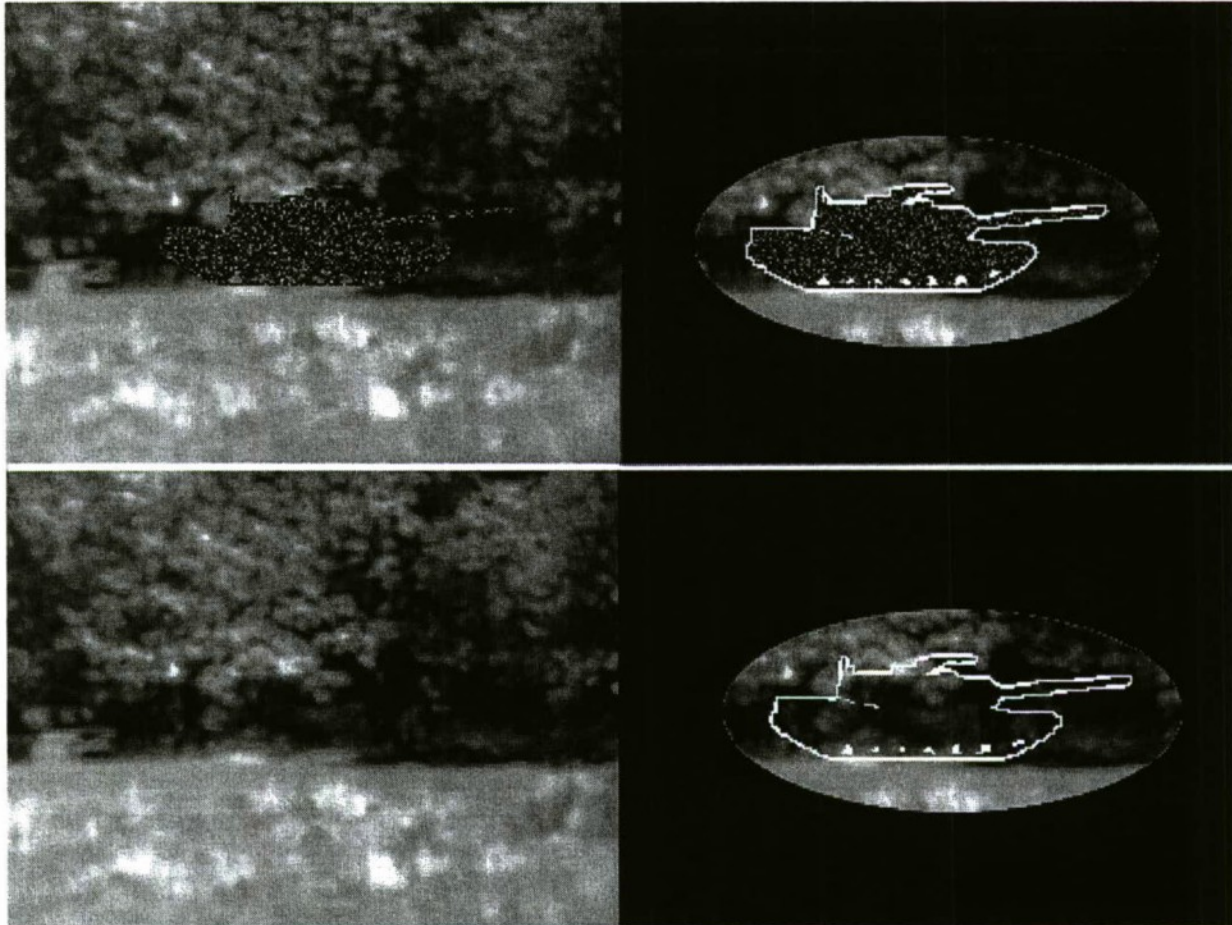
Figure 1 – Equal Signatures? Each of the pictures on the left shows a tank silhouette in a near-infrared scene (outlined to the right). The tanks have the same pixel intensity histograms and will give the same value for most signature metrics. Yet, psychophysically, these pictures are not equivalent.

the computational vision models can really be considered complete, calibrated, and fully validated, although some of these models are validated for specific applications.

While we don't yet have models that can accurately and reliably predict detection probabilities throughout the range represented by the two images in figure 1, there are reliable scaling methods that can help to provide the correct signature level figures-of-merit in a wide variety of situations, including those depicted in this illustration. These scaling methods can provide the psychophysical values with which modeled quantities must correlate. One such method is the Law of Comparative Judgement (LCJ). The LCJ permits us to assign a one-dimensional scale to complex phenomena such as target signature levels even though they may lack an easily identified set of physical attributes and may frequently be a matter of opinion.

LAW OF COMPARATIVE JUDGEMENT

Between 1925 and 1932, Louis Thurstone published 24 articles and a book on how to construct good measurement scales. Today the name Thurstone is synonymous with scaling methods that result in equal-appearing intervals. One of his contributions to the field of psychology is the law of comparative judgement (LCJ).

In the beginning, the LCJ was a psychophysical tool for determining discrimination thresholds and psychological equivalents of physically measurable stimuli. For example, a subject could be presented with a tone of a particular pitch, loudness, and duration, followed by a second tone of the same pitch and duration but not the same loudness. The subject could then be asked whether the second tone was louder or softer than the first one. In this way, investigators could find out how sound pressure translates into perceived sensations. However, the LCJ provides only indirect scaling. As direct means were devised for measuring the same phenomena, psychophysicists turned to these direct methods and the role of the LCJ was gradually eroded. However, abstract sensations (attitudes, opinions, and aesthetic values) provided no physically measurable qualities. Finally, the LCJ came to be primarily a means of characterizing abstract stimuli[1].
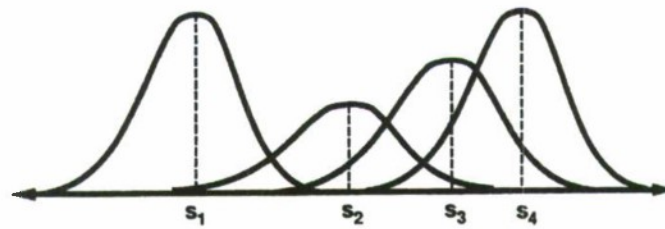
Figure 2 – A Conceptual Psychophysical Scale. This drawing shows 4 stimuli on a hypothetical psychophysical continuum. The horizontal axis indicates the amount of an attribute (e.g. beauty) that each stimulus possesses. The vertical axis indicates the probability that the stimulus will be judged to lie at that point on the continuum at any given time. The regions where the areas under the curves overlap indicate possible inversions.

The fundamental assumption of the LCJ is that when a person is presented with a physical stimulus, it elicits a psychophysical response, and that for any given stimulus, the response may vary from time to time and from individual to individual. Figure 2 shows a conceptual scale on which four stimuli (S_1 to S_4) have been rated. For each stimulus, there is a distribution of responses, which has been assumed to be Gaussian. When the psychophysical values of two stimuli are sufficiently close together, their distributions will overlap as shown in the figure. Under such conditions, it will happen that, for example, S_1 will sometimes be judged greater than S_2 on the psychophysical scale, even though it is actually less. This is called an inversion. It is important to remember that inversions are not "errors" in the normal sense, but the result of random fluctuations in the relationship between physical stimuli and psychophysical responses. In the extreme, two stimuli may be so similar that people cannot distinguish one from the other. In such a case, we would expect that in a forced choice situation, people would be approximately equally likely to pick each of the stimuli and the probability of an inversion would be approximately 0.5.

The LCJ is applied to data from paired comparisons in which people are asked to choose the stimulus that has the greatest (or least) amount of some attribute. For example, tones can be presented in pairs and the subjects could be asked which is loudest (or softest), higher (or lower) in pitch, shorter (or longer) in duration. Pictures can be presented in pairs and the subject can be asked to choose the one that is most beautiful, most relaxing, most representative of a place they would like to be, and so on. Samples of handwriting can be presented in pairs and the subjects can choose the one that is the most readable.

There are many means of ranking stimuli. However, for any given pair of stimuli, the LCJ permits one to do much more than determine which stimulus has most of the attribute being judged. From the amount of overlap in the distributions (represented by the probability of an inversion) one can calculate the distance between the true psychophysical values, provided the stimuli are close enough together that inversions are not too rare. Thus, inversions are a necessary feature of LCJ data, without which numeric scales cannot be ascertained.

As indicated above, people could be given many different tasks for the same set of images. If people were asked to choose the picture that represented the place they would most like to be, we would expect to get substantially different results than if we asked them to pick the one that was the most depressing. Thus, instructions given to the subjects define a task to be performed and greatly affect the choices that are made. Similarly, if we ask our subjects to listen to two tones and choose the one that is higher in pitch, even rudimentary musical training could substantially change the results. Clearly, then, the training and instructions given the subjects can greatly affect the outcome of an LCJ assessment and must be carefully controlled.

USING THE LCJ: A DEMONSTRATION

Procedure

We shall now demonstrate the use of the LCJ by applying it to a practical problem. This demonstration uses a set of 9 images from the Search_2 database[2]. The file names of the images and some of their statistics are shown in table 1. These particular images were chosen because they represented a wide range of signature levels as indicated by mean search time, because they represented a small subset of the targets (all being T-72, M-3, or M-60), and because they represented a broad spectrum of probability of detection. As will be discussed later, it was necessary to keep the set of selected images small.

The images, which had been stored on a CD-ROM in photo-CD format, were read into Adobe PhotoShop® at resolution 5 (3072 x 2048 pixels) and printed 10.24 x 6.827 inches (26.01 x 17.34 centimeters) on 8.5 x 11.0 inch white bond paper using a Hewlett-Packard color LaserJet® 4500N printer.

The subjects (observers) were 13 engineers, scientists, and technicians who work with such images regularly in the

Table I. – Statistics for Selected Search_2 images

Image	Search Time		Nat. Log of Search Time		Visual Lobe		Correct Responses	Search Difficulty (LCJ)
	Arith. Mean	Geom. Mean	Arith. Mean	Geom. Mean	Detect	Identify		
Img0001	14.6	10.1	2.6810	2.3125	0.84	0.06	52	1.6480
Img0013	3.7	3.1	1.3083	1.1314	1.72	1.16	62	0.0000
Img0015	12.4	9.6	2.5177	2.2618	0.29	0.14	36	2.1964
Img0021	15.1	10.9	2.7147	2.3888	1.71	0.29	48	1.3143
Img0022	25.6	21.6	3.2426	3.0727	0.31	0.09	40	2.0914
Img0031	3.5	3.1	1.2528	1.1314	1.65	1.08	62	0.0000
Img0039	34.9	31.6	3.5525	3.4532	0.14	0.07	9	2.4224
Img0042	5.8	4.9	1.7579	1.5892	0.35	0.35	62	0.4920
Img0044	10.6	7.6	2.3609	2.0281	0.27	0.27	57	1.2000
R =	0.848	0.801	0.934	0.930	0.673	0.883	0.842	1.000
R^2 =	0.719	0.641	0.889	0.865	0.453	0.780	0.710	1.000

context of search and target acquisition modeling and psychophysical evaluation. Prior to giving the images to a subject, the images were sorted into order by image number as indicated in table I. Each subject was told to re-sort the images into order from the one in which the target was easiest to find to the one in which the target was hardest to find. The subjects were not immediately told where the targets were in the images, but they were told that information was available when they wanted it. The results of their sorting are shown in table II.

As previously mentioned, LCJ analysis is performed on data from paired comparisons. Furthermore, it is necessary that every stimulus be compared to every other stimulus. Thus, for n stimuli, the total number of comparisons is

$$_nC_2 = \frac{n(n-1)}{2}. \qquad (1)$$

Since this number grows much more quickly than n, it is necessary to keep the number of stimuli in any measurement block relatively small to avoid fatigue among the subjects and to keep the quality of their responses high. At the same time, since inversions are necessary, it is important that stimuli not be too far apart on the psychophysical continuum. While it is possible to obtain meaningful results with as few as 5 well-chosen stimuli, most practical applications limit the number of stimuli to somewhere between 10 and 25.

It was assumed that the subjects' judgements in a paired comparison evaluation would have been entirely consistent with their image collation order. Thus, it was assumed that

any image in the sorted set would have been judged more difficult than any preceding image and less difficult than any later image in the set. This assumption was made because it is statistically most likely, even though inversions (inconsistencies) are common in practice. On this basis, each subject's ordering of the images was converted to a matrix in which a 1 in the i–th row and the j–th column meant that the i–th image was judged easier than the j–th image. Similarly, a 0 meant that the i–th image was judged more difficult than the j–th image. Table III

Table II. – Image Collation Order (Raw Data)

Person	Easiest								Hardest
DT	31	13	42	21	44	15	22	1	39
JeO	31	13	42	1	21	44	22	15	39
BB	31	21	44	42	13	22	5	1	39
DB	1	31	13	42	44	39	22	21	15
DW	31	42	13	21	44	1	22	39	15
JP	31	13	21	42	22	44	39	15	1
GO	31	44	21	13	42	1	15	39	22
JnO	31	42	13	21	44	1	9	22	15
KU	31	13	42	21	1	44	15	22	39
RD	31	13	42	44	15	22	21	1	39
MT	31	13	42	15	21	44	22	39	1
JK	31	13	1	44	42	39	15	22	21
MF	31	13	42	1	44	39	22	21	15

| Table III. —Tally Matrix for subject DT. |
1 = first image preferred.   0 = second image preferred.

		Second Image								
		1	13	15	21	22	31	39	42	44
First Image	1	0	0	0	0	0	0	1	0	0
	13	1	0	1	1	1	0	1	1	1
	15	1	0	0	0	1	0	1	0	0
	21	1	0	1	0	1	0	1	0	1
	22	1	0	0	0	0	0	1	0	0
	31	1	1	1	1	1	0	1	1	1
	39	0	0	0	0	0	0	0	0	0
	42	1	0	1	1	1	0	1	0	1
	44	1	0	1	0	1	0	1	0	0

Table IV. —Tally Matrix from 9 images sorted by 13 people (Frequency of preferring first image).

		Second Image								
		1	13	15	21	22	31	39	42	44
First Image	1	0	1	8	4	8	1	11	2	5
	13	12	0	13	11	13	0	13	10	11
	15	5	0	0	3	6	0	7	0	1
	21	9	2	10	0	9	0	10	3	8
	22	5	0	7	4	0	0	8	0	1
	31	12	13	13	13	13	0	13	13	13
	39	2	0	6	3	5	0	0	0	0
	42	11	3	13	10	13	0	13	0	10
	44	8	2	12	5	12	0	13	3	0

illustrates this process, showing the matrix for the first subject listed in table II.

The matrices for all of the subjects were added, yielding the matrix in table IV.   This matrix was the input to a computer program that applies the LCJ algorithms and produces scale values[3]. For the purposes of this paper, the program will be considered a "black box" with the details of the algorithms considered to be beyond the scope of the present discussion. The interested reader may wish to refer to Copeland and Trivedi[4], Torgerson[5] or Gescheider[6], or contact the author of this paper.

Results

LCJ Search Difficulty

The "search difficulty" values were calculated as described above and are included in the last column of table I. The last two lines of this table show the correlation (r and r^2) between the independent variable (LCJ "search difficulty") and the various dependent variables (metrics) that have been selected. One will observe that the search difficulty correlates very well with several of the metrics, particularly with the natural logarithm of the mean search time (either geometric or arithmetic mean). It seems appropriate to point out that scatter plots generally show very systematic relationships between the search difficulty and most of the selected metrics. However, some of the relationships are decidedly non-linear, causing systematic error when fit to straight lines.   Thus, we find a substantially higher correlation between the search difficulty and the logarithm of the arithmetic mean search time (r = 0.934) than between search difficulty and the mean search time itself (r = 0.848). In the same way, the relationship between search difficulty and probability of detection is also non-linear (see figure 4).

While table I does not have columns for the square and the cube of correct responses, the correlation coefficients are r = 0.923 for the square and r = 0.954 for the cube when compared to the search difficulty (LCJ).   The graph in figure 3 shows the effect of search difficulty (as measured in this LCJ evaluation) on the logarithm of search time. This graph appears to be linear because the vertical scale is logarithmic. The graph in figure 4 shows the effect of search difficulty on the number of correct responses. The trend line shown is a quadratic function with r = 0.939

Goodness of Fit

Testing the goodness of fit between the original data and the LCJ scale values (in this case, search difficulty) is a six-step process.   One must first create a matrix D in which the diagonal elements are zero and for each off-diagonal element,

$$d_{i,j} = S_i - S_j \qquad (2)$$

where $d_{i,j}$ is the element in row i and column j, S_i is the scale
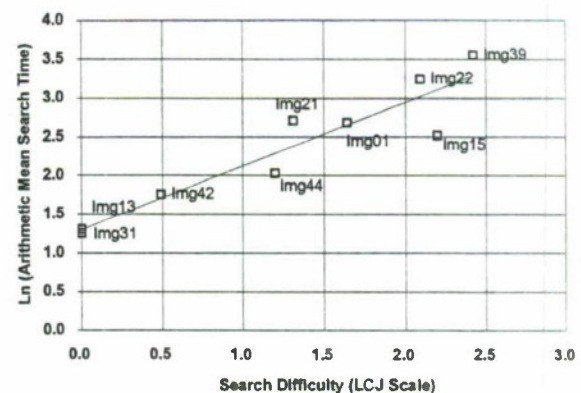
Figure 3 – **Effect of search difficulty on search time.** Nine images from the Search_2 database (r = 0.934)
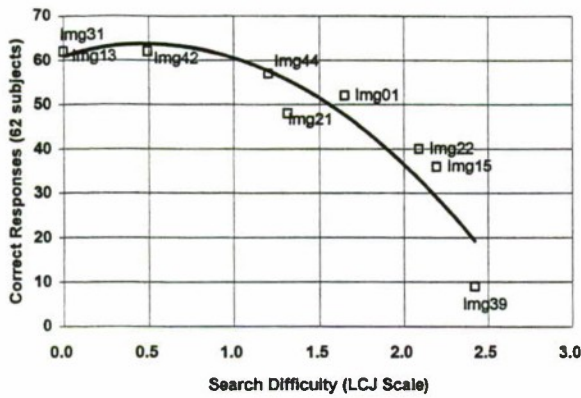
Figure 4 – Effect of search difficulty on correct responses. Nine images from the Search_2 database (r = 0.939)

value for stimulus i, and S_j is the scale value for stimulus j. Because the LCJ scale values produced above were chosen to use one unit normal standard deviation as the scale units, $d_{i,j}$ is the unit normal standard deviate for the separation of the stimulus mean response values. For example, since S_1 = 1.6480 (the scale value for Img0001) and S_3 = 2.1964 (the scale value for Img0015), then $d_{1,3}$ = -0.5484 and $d_{3,1}$ = +0.5484.

The second step is to produce a matrix Z in which each element $z_{i,j}$ is the predicted probability of choosing stimulus i over stimulus j. These probabilities are obtained either from statistical tables or by calculating

$$z_{i,j} = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{d_{i,j}} e^{-x^2/2}. \qquad (3)$$

Next, we calculate expected frequency of occurrence for choosing each stimulus i in preference to every other stimulus j. The elements of this matrix (E) are found by

$$e_{i,j} = Round\,(\,n_{i,j}\,z_{i,j}\,) \qquad (4)$$

where $n_{i,j}$ is the total number of times stimulus i is paired with stimulus j for all observers. Normally, this number is the same for all stimuli, in which case all $n_{i,j}$ can simply be replaced by n. For our example, the expected frequency of occurrence is given in table V.

The fourth step is to calculate

$$\chi^2 = \sum_{i=1}^{m} \frac{(o_{i,j} - e_{i,j})^2}{e_{i,j}} \qquad (5)$$

where the values $o_{i,j}$ are the observed frequencies of occurrence from table IV. The upper limit of the summation is

$$m = \frac{k(k-1)}{2} \qquad (6)$$

where k is the number of stimuli in the experiment. However, the number of elements in the matrices O and E is k^2, and we are not using all of them, so it is necessary to define the selection process. In this case, we will select $o_{i,j}$ and $e_{i,j}$ only if $z_{i,j} \geq 0.5$. Furthermore, when $z_{i,j} = 0.5$, then $z_{j,i}$ is also 0.5 and $o_{i,j} - e_{i,j} = o_{j,i} - e_{j,i} = 0$. In these cases, we will use either of these differences, but not both. For our example, χ^2 = 16.3335.

We shall next calculate ν, the degrees of freedom as

$$\nu = m - k \qquad (7)$$

where m comes from equation 6 and k is again the number of stimuli. In the example, ν = 27.

Finally, the goodness of fit is determined by integrating the chi-squared distribution from 0 to χ^2 with ν degrees of freedom to obtain the probability of error. (The confidence is 1 minus the probability of error.) Normally one would not perform the integration, but use tables instead. However, the most common chi-squared tables in textbooks and most other sources only go up to 30 degrees of freedom. In our current, very limited case, ν = 27. With 10 stimuli, the degrees of freedom increase to 35, and with 25 stimuli, it would be 275. It is clear that tables will normally not serve our needs.

There are at least two solutions to this dilemma. Available computer software can be used to calculate the probabilities. If you lack such software, the NCSS Probability Calculator[7] should serve your needs and is available free over the internet. Also if ν > 30, the formula

$$d = \sqrt{2\chi^2} - \sqrt{2\nu - 1} \qquad (8)$$

may be used to calculate the normal standard deviate d associated with χ^2 and ν[8]. You may then refer to widely available tables for probabilities associated with the normal (Gaussian) probability density function. Such tables are found in statistics textbooks and standard mathematical tables. It may be sufficient to refer to table VI, which gives

Table V. – Expected Frequency of preferring first image (13 people).

		Second Image								
		1	13	15	21	22	31	39	42	44
First Image	1	0	1	9	5	9	1	10	2	4
	13	12	0	13	12	13	7	13	9	12
	15	4	0	0	2	6	0	8	1	2
	21	8	1	11	0	10	1	11	3	6
	22	4	0	7	3	0	0	8	1	2
	31	12	7	13	12	13	0	13	9	12
	39	3	0	5	2	5	0	0	0	1
	42	11	4	12	10	12	4	13	0	10
	44	9	1	11	7	11	1	12	3	0

Table VI – Probabilities associated with key values of the normal standard deviate.

D	Probability of error	Confidence
-1.282	0.10	0.90
-1.645	0.05	0.95
-2.326	0.01	0.99
-2.576	0.005	0.995
-3.090	0.001	0.999

Source: NCSS Probability Calculator

five key values of d, the probability of an error, and the corresponding confidence levels.

In the case of our example, equation 8 cannot be used because we have only 27 degrees of freedom. The NCSS Probability Calculator gives 0.054 for the probability of error and 0.946 for the confidence.

Repeatability

The group of observers in our demonstration sorted 5 of the images several days prior to the evaluation recorded in table II. The data was processed as described above and search difficulty values were calculated. When the scale values from the two sorting exercises were compared for these 5 images, the slope of the regression line was 0.997 and the correlation coefficient was r = 0.980. This indicates that the results were highly repeatable. However, since the process was not repeated with a different set of subjects, we cannot safely draw any conclusions about the performance of any other group of individuals or the population as a whole.

Discussion

The LCJ evaluation that was outlined above was relatively quick and easy compared to a properly run search experiment. At the same time, it correlates very well with search time and probability of detection. It would appear to be a highly effective tool for determining the relative strength of target signatures. At the same time, the LCJ has certain limitations.

The LCJ search difficulty scale that we obtained above is a psychophysical scale with no natural zero point and units that have no obvious relationship to useful quantities such as average time required to detect the target or probability of detection. Furthermore, the scale will change from one experiment to the next with no common reference. Thus, one might easily ask what advantage there is to such measurements. Is there any reason to use the LCJ in preference to other psychophysical measures or methods? I would like to suggest that there are numerous circumstances that might lead one to use the LCJ either in preference to other methods or in conjunction with them.

First, it is necessary to realize that the lack of a natural zero and a physically meaningful scale are really not significant problems. Detection time and probability of detection, while seemingly more meaningful are actually relative as well. The skill of the observers, the conditions under which the images are viewed, and many other variables in addition to the images themselves, will all affect the detection time and probability of detection. Observers who are more or less skilled, more or less effectively trained and motivated, or who are viewing images of varying quality and magnification will give varying results. Thus, in either case, two things are required: calibration standards and conversion formulas.

For example, in the case of the nine stimuli in the exercise above, the conversion from search difficulty to arithmetic mean search time in seconds can be expressed as

$$t \approx 5.21\,e^{0.82s} + 4.78. \qquad (9)$$

where t is time in seconds, s is the search difficulty, and e is the base of the natural logarithms. However, one must bear in mind that this formula applies only to the relationship between the search difficulty as measured by the data from the 13 Night Vision employees and search times for the 62 observers in the TNO test. It is likely that the 13 Night Vision employees could predict the search time on other images in the Search_2 set. It may also be that search times from the Search_2 data could be used to predict the search difficulty for other images. However, he who would extend this relationship to search difficulty values for other images sorted by other people or to search times in other search experiments would be making a potentially serious error.

Even so, all is not lost. Just as there was a day when two marks were scribed on a platinum-iridium bar to define a meter, other standards of measurement have been defined before and since. In the same way, useful perceptual standards can also be defined. However, rather than continue with search time and search difficulty, let us examine another phenomenon – visual clutter.

USING THE LCJ: A CLUTTER SCALE

The LCJ is primarily a tool for building measurement scales. Thus, we examine clutter as an example of an important quantity for which we have no accepted scale. Our purpose is to see how the LCJ could be used to build a standard reference scale. This relates to our primary topic of target signature evaluation in that target signatures must be evaluated in the context of a background and clutter is one of the most fundamental ways of characterizing backgrounds.

Definition of Visual Clutter

Clutter has been defined as "scene elements similar enough in size and contrast to the [target] that each one has to be

considered in detail as a potential target"[9]. The concept of clutter is pervasive and generally describes distracting, annoying, and unwanted signals or returns when any of a wide variety of sensors is used. It is often discussed but seldom precisely defined. We shall use the phrase "visual clutter" in this paper to apply to any situation where there is a person using their eyes to examine a scene in which there is clutter, whether they are using "bare eyes" or an imaging sensor.

For many years, investigators have known that an observer's performance depends on many factors, including clutter. Schmieder[10] has probably been more influential than anyone else in the quest to subject visual clutter to quantification and analysis, but the proliferation of clutter metrics is testimony to the fact that none of the metrics are convincingly successful. However, the LCJ could be used as a tool in establishing a clutter scale that would be perceptually meaningful, extensible, and widely applicable.

Establishing a Unit of Visual Clutter

The first step in establishing a perceptual image clutter scale would be to select a set of images exhibiting a wide range of clutter levels. In order to maintain generality, they should represent numerous locales and clutter types. Since many feel that clutter must be understood in the context of the target, the set should include images with targets as well as images without targets. Initially, it would probably be satisfactory to have only military ground vehicles as targets.

From the initial set of images, a training package should be prepared so that observers can be taught what clutter is and so that they can become familiar with the size scales of the images in the set. This will help to make results repeatable, a necessary feature. A set of test stimuli would also need to be selected and should be distinct from the training set.

A pool of observers would also be required. The pool would need to be large enough that aberrant results from any one observer would have negligible effect on the results. Experience has indicated that at least 25 observers would be desirable. The observers would first be trained using the training set along with appropriate commentary. When they were fully trained, they would participate in a paired comparison evaluation of the test images. Their task would be to choose the image in each pair that had the most (or least) visual clutter.

When all observers had completed the paired comparison evaluation of the test set, LCJ statistical analysis would be used to obtain the perceptual image clutter scale. At this point, the scale would be arbitrary. Probably the image that had the lowest clutter would be selected as the zero point.

From the test images, a subset would be selected as a reference set. Images that had the same, or nearly the same

perceptual image clutter values would be culled. An attempt would be made to select a relatively small number of images that spanned the entire scale, and were evenly distributed between the extremes, but with no gaps. It would be best if about 1 unit normal standard deviation separated the individual images in the reference set. Probably 1 unit normal standard deviation would be selected as the scale unit.

It would be highly desirable to repeat the evaluation with a second pool of observers in order to establish whether or not the scale is indicative of a broader population. Actually, several replications would be ideal. If this could be done, the first replication should be with a group as similar to the first one as possible. Thereafter, greater liberties could be taken with the makeup of the observer pool in order to observe how robust the scale actually was.

Evaluating Clutter Levels

Having established a clutter scale for one set of images, one would naturally want to determine where other images were on the same scale. This could be done in any of at least 3 ways.

Quick Estimate

For a quick estimate of the clutter level in any image, anyone who was well versed in the perceptual image clutter scale could simply compare a new image to the reference images. Assuming there was nothing unusual about the image, they would be able to tell where it belonged on the scale, probably within about half of a unit. Tests of this method could be verified by one of the other methods to determine reliability.

LCJ Method

A second method of determining the clutter level in one or more images would be to mix new images with some or all of the reference set and perform a paired-comparison LCJ evaluation as described above. The results from the new evaluation would be used in conjunction with a linear transformation of scale values that would minimize the error for the reference images. This linear transformation could be determined by simply doing a linear regression between the standard values for the reference images and the values obtained for them in the new evaluation. The correlation coefficient obtained would be a measure of the reliability of the values assigned to the new images.

Jury Method

A third method would be to have a panel of "experts" who were all familiar with the perceptual image clutter scale assign clutter values to each of the new images. This would be more reliable and precise than the quick method above at the same time that it would be quicker and easier than the

LCJ method. The major drawback to this method would be that there would be no ready means of evaluating the reliability of the values assigned to the new images.

Extending the Scale

If this methodology were employed, we might in time encounter clutter levels that were beyond the limits of the original set. There is nothing about the methodology described above that limits it to interpolation alone. In time, more reference images could be added to the set by the LCJ method outlined above. The only requirement in extrapolating beyond the original set is that no new set of images can be added if any continuous subset lies more than about one standard deviation beyond either end of the scale (depending on the number of observers in the pool). However, in such a case, selection of enough images with a variety of intermediate clutter levels should provide the necessary continuum.

Observer Pools and the Population

We previously alluded to the fact that different populations might give different results. If this methodology were adopted for establishing a clutter scale, it would be wise to determine how stable the results were across these various populations. For example, it might be that trained military personnel would not give the same results as civilian clerical employees. On the other hand, since we are only asking individuals to make relative judgements ("Which image has the most clutter?") as opposed to quantitative judgements ("How much clutter does this image have?"), we may find that the numbers obtained are quite stable over a broad spectrum of the human population. If the latter were true, it would be fortunate and knowing that it was true would permit various economies since trained military personnel are not always readily available at research facilities. At the same time, this cannot be assumed.

Analytical Methods

Naturally we would prefer to have analytical means of determining clutter levels rather than rely on psychophysical measures. However, we must remember that the human eye-brain system is most often the standard against which performance is rated. Having a reliable scale would be of great value in testing analytical methods because investigators would know what the "correct answers" are. Even if analytical methods were only able to tell which reference image a new image was most like, that would be a step in the right direction and eventually it could eliminate the need for paired comparisons and juries.

CONCLUSIONS

The Law of Comparative Judgement (LCJ) has great potential for helping us evaluate target signatures. There are two ways in which this potential might be realized. First, the LCJ can provide relatively quick, easy answers to questions that involve a complex set of variables such as we encounter when evaluating target signatures. It has been shown, for example, that the LCJ can give good estimates of mean search time using a methodology that is much quicker and easier that a traditional search experiment. When relative answers such as "Which is better?" and "How much better is it?" will suffice, or when there is a known relationship between LCJ scale values and important measures of effectiveness, the LCJ can be a highly effective tool. The LCJ can also be used to build scales for qualities that are difficult to quantify. This is perhaps where its greatest potential lies. To explain how this works, a scheme has been outlined for creating a perceptual image clutter scale. Such a scale could provide important benchmarks in an area of image understanding that has long been in need of an anchor. Both of these applications could contribute greatly to the important area of target signature evaluation, search, and target acquisition.

REFERENCES

[1] Green, D.M., and Swets, J.A., *Signal Detection Theory*, Peninsula Publishing, Los Altos, CA, 1988.

[2] Toet, A., Bijl, P., Kooi, F.L., and Valeton, J.M., *A high resolution image data set for testing search and detection models, (Report TM-98-A020)*, TNO Human Factors Research Institute, Soesterberg, The Netherlands, 1998.

[3] Copeland, A.C., *Xpet_pairs_LCJ* (Computer Program), Contract DAAK-70-93-C-0037, U.S. Army, CERDEC, Fort Belvoir, VA, 1997.

[4] Copeland, A.C., Trivedi, M.M, and Ravichandran, G., *Developing a Quantitative Basis for Synthesis, Analysis, and Assessment of Complex Camouflage Patterns, Contract DAAK-70-93-C-0037*, U.S. Army, CERDEC, Fort Belvoir, VA, 1997.

[5] Torgerson, W.S., *Theory and Methods of Scaling*, Krieger Publishing, Malabar, FL, 1985.

[6] Gescheider, G.A., *Psychophysics: method, theory, and application*, Erlbaum Assoc. Publishers, Hillsdale, NJ, 1985.

[7] *NCSS Probability Calculator* (Computer Program), NCSS Statistical Software, Kaysville, UT, 1995.

[8] Weast, R.C, ed., *C.R.C. Standard Mathematical Tables, 13th Ed*, The Chemical Rubber Co., Cleveland, OH, 1964.

[9] Lloyd, J.M., "Fundamentals of Electro-Optical Imaging Systems Analysis" *in The Infrared & Electro-Optical Systems Handbook, Volume 4: Electro-Optical Systems Design, Analysis, and Testing*, M.C. Dudzik, ed., SPIE, 1993.

[10] Schmieder, D.E. and Weathersby, M.R., "Detection Performance in Clutter with Variable Resolution," *IEEE Transactions on Aerospace and Electronic Systems*, 19:4, 1983.

Fundamentals of the Human Foveal Vision System

G.R. Gerhart, R.M. Matchko and R. Goetz
Vehicle Detectability Research Team
AMSTA-TR-R / MS 263
Warren, MI 48397-5000
gerhartg@tacom.army.mil

ABSTRACT

This paper presents a mathematical model, which includes a closed-form equation (the ABCt law) that generates the entire Blackwell-McCready (BM) threshold data set. It derives relationships among the four fundamental parameters of foveal vision: target area **A**, background luminance **B**, threshold contrast **C** and stimulus presentation time t. It shows that graphs of log (t) as a function of log (ABC) are hyperbolas whose asymptotic regions are related to the well-known laws of Weber-Fechner and Bloch. It unifies important relationships associated with target/background scene parameters as they relate to the human foveal vision process. The constants associated with the simple empirical laws of Ricco, Blackwell, Weber-Fechner and Bloch can easily be obtained for a large range of target /background conditions and stimulus presentation times. Conditions for the most efficient stimulation of the human visual system are quantified and expressed in terms of the total energy for specific detection tasks.

1. INTRODUCTION

Along the visual axis, intersecting the retina, there is a small area on the retina, the fovea, which provides the greatest visual acuity and spatial resolution. Covering about one degree of visual angle in diameter, the fovea contains the greatest concentration of cones, but no rods. The BM data set relates to this part of the human vision system.

In an earlier paper[1,2], a hyperbolic curve fitting algorithm was derived for the BM data set. It accurately generates the threshold contrasts (**C**) of the circular targets for a wide range of target areas (**A**) and uniform background ambient luminance (**B**) for a given stimulus duration time (**t**).

Typical displays of the BM data show plots of log C as a function of log (A) or log (B) with the presentation time t held constant, as shown in Figures 1 and 2 for a presentation time t = 1/3 second. The circular markers represent points from the BM data set and the solid lines were generated from the hyperbolic curve fitting algorithm[1,2].
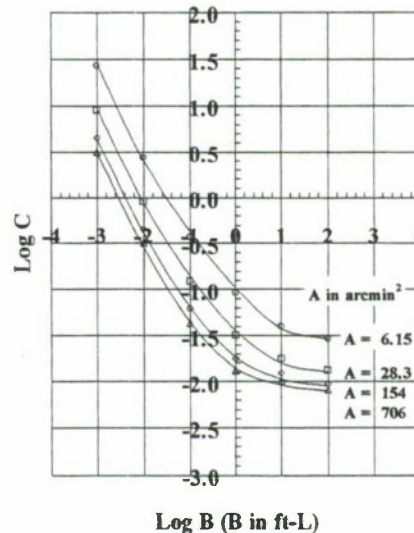
Figure 1. Log C versus Log A for t = 1/3 second.
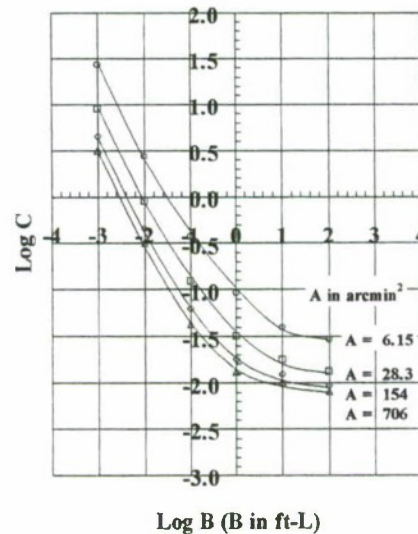
Figure 2. Log C versus Log B for t = 1/3 second.

Previous hyperbolic curve fitting algorithms were developed for all three 2-D projections of the BM data set with constant presentation time t. Each geometric projection exhibited a number of interesting characteristic features of the human visual system. A significant accomplishment in the previous work was the recognition that the Ricco and Blackwell laws can be combined to give what we referred to as the ABC law for foveal vision. This paradigm stated that the product of A, B and C was a constant for small target areas (A< 6 arcmin^2) and low ambient light levels (B < 0.1 ft-L). Figure 3 is a graph of the ABC constant as a function of log (t). It was also shown that the largest variations in human performance, as a function of B, occurred during the time interval before dawn and after dusk. A subsequent paper[3] extrapolated these results through a redefinition of ambient luminance to real world scenarios.

Our continued efforts to characterize and further analyze the BM contrast discrimination data set has led to the discovery of a closed-form equation (the ABCt law) which generates the entire BM smoothed threshold data set. In particular, it shows that plots of log t versus log ABC are hyperbolas (one branch) with one vertical asymptote and the other with a slope of -1, as shown in Figure 4. The asymptotic regions are related to the well-known laws of Weber-Fechner and Bloch. The ABCt law unifies and quantifies many of the important interrelationships associated with the physical environment as they relate to the foveal vision process. It unifies under a single formalism the empirical laws of Ricco, Blackwell, Weber-Fechner and Bloch. The associated constants with these four laws are obtained for a wide range of target/background parameters and stimulus presentation times. A dimensional analysis of the product ABCt shows it corresponds to the minimum energy required to detect the target. This formalism is useful for predictive models and provides quantitative estimates of human performance parameters over a wide range of conditions. It should be relevant to many applications of tactical interest and serve as a performance baseline for more advanced computational vison models. The integration of the effects of light energy over stimulus area (spatial summation) and over time (temporal summation) can also be quickly classified and quantified.

The Early Empirical Laws

A number of simple empirical laws have been formulated to describe human visual system performance. Ricco=s law (CA = constant) is relevant for small targets with large relative contrast values. It results from the fact that diffraction effects dominate for object sizes less than some critical angle.

Blackwell=s law (CB = constant) is applicable for low ambient
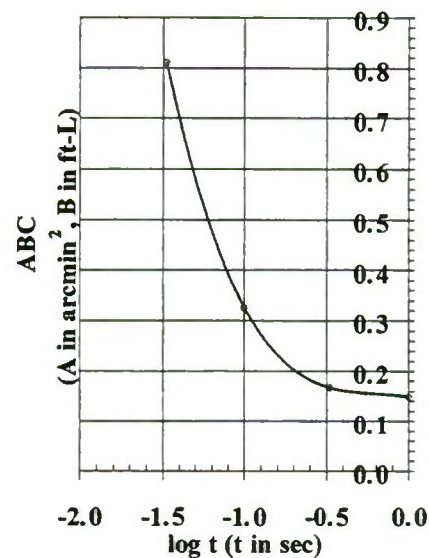
Figure 3. The ABC constant versus log t for B <0.1 ft-L and A < 6 arcmin^2.
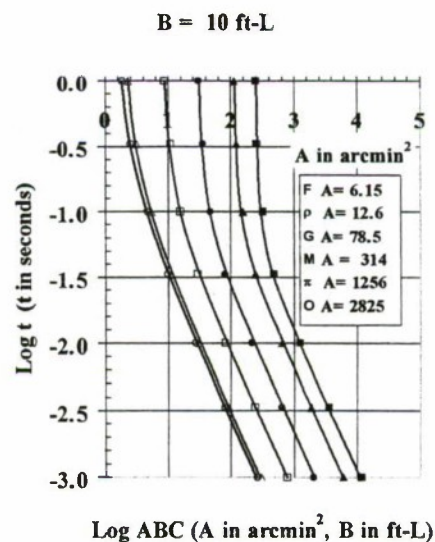
Figure 4. Log t versus Log ABC for B = 10 ft-L.

luminance levels, and it illustrates the high, relative sensitivity of the foveal vision system during ambient lighting conditions, which typically occur before dawn and after dusk. The

Weber-Fechner law (C = constant) applies for large, highly resolved targets when observer thresholds are nearly constant and independent of size or range. This phenomenon is indicative of noise limiting performance of the human observer and relatively independent of the target characteristics A and C. Bloch=s law, which applies for short stimulus presentation times, states that the product of intensity and time (ABCt) is a constant for t less than some critical value, and it is proportional to some photochemical effect in the human vision system.

In each of these simple paradigms the product of two parameters is constant for some range of asymptotic conditions. Although the ABCt law provides a means to obtain these constants, it can also provide numerical values for describing the entire range of target/background variability affecting human foveal vision.

Units

Although target areas have been expressed in arcmin^2, they can easily be converted to m^2 with the knowledge that the BM observers were 10.25 feet from the target. Background luminance levels are expressed in foot-Lamberts with 1 ft-L = 3.426 cd/m^2. Threshold contrast is the ratio of the difference between target luminance and background luminance to the background luminance that elicits a response from the vision system. Therefore, the product ABC has units of light intensity. If target area is expressed in m^2 and luminance in cd/m^2, the units for the product ABC are candelas (cd). For a point source, 1 cd = 1 lumen per steradian. Since the lumen is a unit of light power, the product ABCt represents the minimum energy required by the foveal vision system to detect the BM targets.

2. BACKGROUND INFORMATION

Comprehensive experimental data sets for human observers which relate scene luminance, visual contrast, target dimension and stimulus presentation time are relatively few in number. Blackwell published several experimental studies[4,5] beginning in the World War II time period. His earliest work[4], generally known as the Tiffany data set, was intended for use in military applications with less than maximum visibility or relative target/background contrast. The Blackwell-McCready[5] (BM) data set attempted to provide the first comprehensive body of data in which background luminance level, target size and duration were all studied over wide ranges of practical interest. The data was collected using procedures that allow the analyst to convert the threshold data to various probabilities of detection levels. The Tiffany data consisted of three basic studies conducted over a comprehensive range of background luminance and target size. The first and second perception experiments were forced-choice detection tasks, which required observers to search for circular targets in a uniform background over a six-second-presentation time. The third experiment used bright targets in known positions for presentation times necessary to obtain maximum detection probabilities.

The basic 1958 BM data set was collected from two Ahighly motivated and experienced observers@. The original data consists of 81,000 observations in 162 experimental sessions. Target diameters ranged from 0.802 to 51.2 minutes of arc. Background luminance values varied from 0.001 to 100 ft-L which corresponds to the range of the average pupil diameter of the human eye (2.5 mm to 7.5 mm)[6]. Unlike the Tiffany data sets, presentation times were varied from 0.001 to 1 second and more attention was given to resolving fixation concerns. The BM targets always appeared in a known location at the center of the screen where the observers initially fixated their foveal vision. Associated with each value of A, B and t was a contrast threshold value C corresponding to a detection probability of 0.5. After a smoothing process, the smoothed threshold data consisted of 546 threshold contrasts corresponding to 13 target areas, 6 background luminance values and 7 presentation times. Since the product of ABC does not change for target areas less than six arcmin^2, 126 of the BM threshold contrasts are redundant. The ABCt law refers to the remaining 420 threshold contrasts.

3. HYPERBOLIC CURVE ALGORITHM

Basically, the ABCt law is the equation of a family of hyperbolas. Its conception was prompted from an analysis of plots like those shown in Figure 4. Each plot is for a specific value of A and B. The markers represent values from the BM data set and the solid curves were generated from the ABCt law. A few fundamentals involved in the derivation of the ABCt law are given below. Equation 1 is the simplest mathematical form for a hyperbola with foci along the y-axis, semi-axes a and b and center at the origin of the (x, y) coordinate system.

$$\frac{y^2}{a^2} - \frac{x^2}{b^2} = 1 \tag{1}$$

The plots in Figure 4, however, appear to be hyperbolas described by Eq. (1) in a coordinate system which has been translated and rotated. This is illustrated in Figure 5 which shows the 2-D, geometric transformation between the (x, y) and (x', y') coordinate systems. Equations (2) and (3)

represent a counterclockwise rotation (θ) and a translation of the origin to the point (h, k)

$$x' = (x-h) \cos \theta + (y-k) \sin \theta \qquad (2)$$
$$y' = (y-k) \cos \theta - (x-h) \sin \theta \qquad (3)$$

Equation 1 is a special case of a general equation of second degree:

$$A x^2 + B x y + C y^2 + D x + E y + F = 0 \qquad (4)$$

where $B^2 - 4AC > 0$. In the (x',y') coordinate system:

$$A' x'^2 + B' x' y' + C' y'^2 + D' x' + E' y' + F' = 0 \qquad (5)$$

where

$$A' = b^2 \sin^2 \theta - a^2 \cos^2 \theta \qquad (6)$$
$$B' = (a^2 + b^2) \sin 2\theta \qquad (7)$$
$$C' = b^2 \cos^2 \theta - a^2 \sin^2 \theta \qquad (8)$$
$$D' = 2 b^2 k \sin \theta - 2 a^2 h \cos \theta \qquad (9)$$
$$E' = 2 b^2 k \cos \theta + 2 a^2 h \sin \theta \qquad (10)$$
$$F' = b^2 k^2 - a^2 h^2 - a^2 b^2 \qquad (11)$$

The asymptotes in the (x=, y=) coordinate system are characterized by:

$$y' = \frac{(a \cos \theta - b \sin \theta) x' + (a h - b k)}{a \sin \theta + b \cos \theta} \qquad (12)$$

$$y' = \frac{(a \cos \theta + b \sin \theta) x' + (a h + b k)}{a \sin \theta - b \cos \theta} \qquad (13)$$

If m_1 and m_2 are the slopes of the asymptotes, then we get

$$\tan 2\theta = \frac{m_1 + m_2}{m_1 m_2 - 1} = \frac{B'}{C' - A'} \qquad (14)$$

4. The ABCt Law

Since plots of log t versus log ABC are hyperbolas (one branch) with one vertical asymptote and the other with a slope of -1, according to Eq. (14), θ = 67.5E and it can be shown that b = a tan (67.5E). Using these two results and Eqs. (5) to (8) yields AN = BN and CN = 0. Hence, an equation that will generate the entire BM smoothed threshold data set is of the form

$$A' (x'^2 + x' y') + D' x' + E' y' + F' = 0$$
$$\{\text{target area } A > 6 \text{ arcmin}^2\} \qquad (15)$$

Log ABC
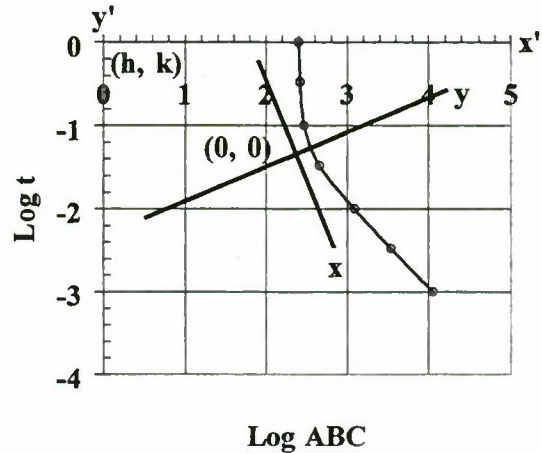
Figure 5. Transformation of coordinates.

where the coefficients are given by Eqs. (6) to (11) and the coordinates (x', y') correspond to (log ABC, log t). The restriction $A > 6$ arcmin^2 is founded on the fact that plots of log t versus log ABC for values below this limit are all redundant. Eq. (15) is the ABCt law. Inspection of Eqs. (6) to (11) shows that the only parameters to identify are a, h and k. An analysis of the entire BM data set shows that each parameter, a, h and k can be expressed as a function of the target area A and the background luminance B. Plots of a, h and k versus log A are also hyperbolas; however, one asymptote is horizontal and the other has a negative slope which depends on the ambient luminance B. A pictorial representation of this scenario is shown in Figure 6. This suggests that the a, h and k parameters can be expressed in the form of Eq. (5), where the coordinates (x', y') now correspond to (log A, a), (log A, h) or (log A, k).

For distinctness, all parameters relating to the (log A, a) hyperbolas will have the subscript A1@, all parameters relating to the (log A, h) hyperbolas will have the subscript A2@ and all parameters relating to the (log A, k) hyperbolas will have the subscript @3A. For example, the parameters A', B', C, D', E', F, a, b, h, k, and θ of Eqs. (5) to (11) correspond to A'₁, B'₁, C'₁, D'₁, E'₁, F'₁, a₁, b₁, h₁, k₁, and θ₁ of the (log A, a) hyperbolas.

Since one asymptote of the a, h and k parameter hyperbolas is

246

horizontal, $b_n = a_n \cot \theta_n$ and $A_n' = 0$. Hence, analogous to Eq. (5)

$$B'_n\, x'\, y' + C'_n\, y'^2 + D'_n\, x' + E'_n\, y' + F'_n = 0 \qquad (16)$$

where, $n = 1, 2, 3$; $x' = \log A$ and $y' = a, h, k$.

Through a comprehensive curve-fitting procedure, the parameters a_n, h_n, k_n and θ_n were adequately described as given below. Using $z = \log B$, the a, h and k parameters are described by

$a_1 = 0.004$, $b_1 = a_1 \cot \theta_1$
$h_1 = 0.016788\, z^5 - 0.006235\, z^4 - 0.118709\, z^3 + 0.149452\, z^2$
  $- 0.042503\, z + 0.898705$
$k_1 = 0.000592\, z^5 - 0.007458\, z^4 - 0.034943\, z3 + 0.037570\, z^2$
  $+ 0.186497\, z + 0.279860$
$\theta_1 = -0.082083\, z^5 - 0.320417\, z^4 + 0.177917\, z^3$
  $+ 1.570417\, z^2 + 0.814167\, z + 180.4200$

$a_2 = 0.020$, $b_2 = a_2 \cot \theta_2$
$h_2 = 0.014778\, z^5 + 0.040227\, z^4 - 0.066438\, z^3 - 0.154833\, z^2$
  $+ 0.048312\, z + 1.175540$
$k_2 = -0.003207\, z^5 - 0.006359\, z^4 + 0.004473\, z^3$
  $- 0.060491\, z^2 - 0.162769\, z - 0.179468$
$\theta_2 = -0.025000\, z^5 - 0.208333\, z^4 - 0.383333\, z^3 + 0.78333\, z^2$
  $+ 2.633333\, z + 192.300$

$a_3 = 0.192$, $b_3 = a_3 \cot \theta_3$
$h_3 = 0.002347\, z^5 + 0.016055\, z^4 + 0.023844\, z^3$
  $- 0.044514\, z^2 - 0.089004\, z + 0.993574$
$k_3 = 0.001136\, z^5 + 0.001543\, z^4 - 0.020823\, z^3$
  $- 0.105677\, z^2 - 0.319653\, z + 1.164966$
$\theta_3 = -0.009167\, z^5 - 0.071667\, z^4 - 0.100833\, z^3$
  $+ 0.441667\, z^2 + 1.240000\, z + 198.800$

$$\qquad (17)$$

Example

The following example uses the ABCt law (Eq. 15) to calculate the threshold contrast for a target area $A = 78.54$ arcmin^2, a background luminance $B = 10$ ft-L and a presentation time of 1/3 second. From Eq. (17):

$a_1 = 0.004$, $h_1 = 0.89750$, $k_1 = 0.46212$, $\theta_1 = 182.58$,
$b_1 = 0.089$

$a_2 = 0.020$, $h_2 = 1.05759$, $k_2 = -0.40782$, $\theta_2 = 195.10$,
$b_2 = 0.074$

$a_3 = 0.192$, $h_3 = 0.90230$, $k_3 = 0.72149$, $\theta_3 = 200.30$,
$b_3 = 0.519$.

$$\qquad (18)$$

From Eqs. (7) to (11):

$B'_1 = 0.00071$, $C'_1 = 0.00786$, $D'_1 = -0.00030$, $E'_1 = -0.00728$,
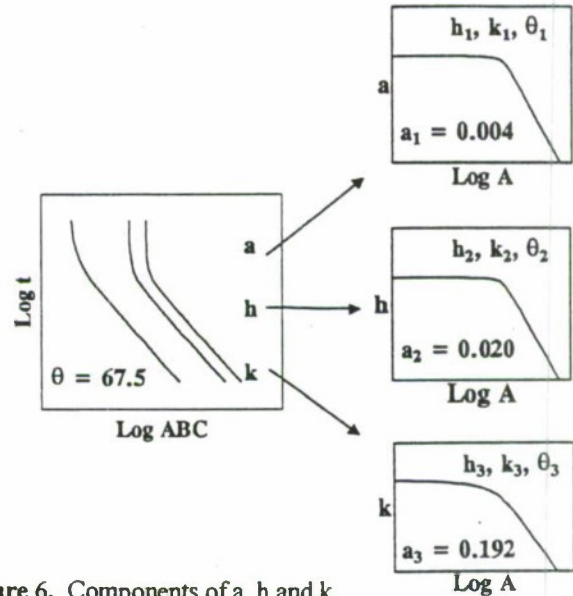
Figure 6. Components of a, h and k.

$F'_1 = 0.00167$;
$B'_2 = 0.00296$, $C'_2 = 0.00509$, $D'_2 = 0.00198$, $E'_2 = 0.00411$,
$F'_2 = 0.00046$;
$B'_3 = 0.1993$, $C'_3 = 0.2325$, $D'_3 = -0.07248$, $E'_3 = -0.3876$,
$F'_3 = 0.1003$.

$$\qquad (19)$$

Inserting the coefficients from Eq. (19) and $\log A = 1.8951$ into Eq. (16) produces

$0.00786\, a^2 - 0.00593\, a + 0.00110 = 0$
$0.00509\, h^2 + 0.00972\, h + 0.00421 = 0$
$0.233\, k^2 - 0.00991\, k - 0.0371 = 0 \qquad (20)$

Using the lower branch of the hyperbola, we now know that $a = 0.329$, $h = -1.25$, $k = -0.378$, $\theta = 67.5E$ and $b = 0.794$.

Applying the above results to Eqs. (6) to (11) yields

$A' = 0.522$, $D' = -0.337$, $E' = -0.432$ and $F' = -0.147 \qquad (21)$

Inserting these coefficients into Eq. (15) with $yN = \log t = -0.4776$ gives

247

$$0.522\ x'^2 - 0.586\ xN + 0.059 = 0 \qquad (22)$$

Solving for xN produces log ABC = 1.01 and C = 0.0130 or 1.3 %. The BM data set gives C = 0.014 or 1.4 %. These results are very typical of solutions obtained from the ABCt law.

Quantifying Ricco=s Law and Blackwell=s Law

Since the ABCt law is a closedBform equation, it can be used to quantify many of the simple empirical laws that have been formulated to describe the performance of the fovea of the human eye. Ricco=s law is illustrated in Figure 2 where the plots for all A < 6 arcmin^2 are coincident and the coefficients in Eq. (15) are constant for a given value of B. Hence, for a given background luminance B and exposure time t in this region, Eq. (15) reduces to a simple quadratic equation with log (AC) as the variable when A < 6 arcmin^2. The value of the product AC is the Ricco constant for the assumed values of B and t. For example, if B = 10 ft-L and t = 1/3 second, Eq. (15) simplifies to 0.8517 (log AC)^2 + 1.863 log AC + 0.8253 = 0 and log (AC) = -0.617 or AC = 0.2415; the original BM data set gives AC = 0.247. This phenomena stems from the fact that below a minimum critical angle, diffraction effects primarily determine the apparent observer image size while the target range and physical dimensions determine the illumination and subsequent apparent threshold contrast. The end result is that the product of contrast threshold and target area is a constant whose value is a function of ambient luminance and presentation time.

Blackwell=s law is illustrated in Figure 3 where plots for B < 10^{-2} ft-L are nearly coincident and the coefficients in Eq. (15) are nearly constant for a given value of A. Hence, in this region, for a given target area A and exposure time t, Eq. (15) reduces to a simple quadratic equation with log (BC) as the variable when B < 10^{-2} ft-L. The value of the product BC is the Blackwell constant for the assumed values of A and t. This region describes a range where rod vision typically dominates the human vision process. Cone vision dominates for brighter ambient luminance conditions, such as direct sunlight, when B >> 0.1 ft-L which produces noise or contrast limiting conditions. The transition between B < 1 ft-L and B >0.01 ft-L contains a mixture of both rod and cone vision such as occurs near dawn or dusk. The ABCt law generalizes Ricco's and Blackwell's laws and gives some particularly simple relationships between the relative target/background contrast, target area and ambient luminance which are useful for military and commercial applications.

Bloch=s Law

Figure 4 reveals that there is an approximate straight-line segment of slope B1 associated with each plot.

Log ABC (A in arcmin^2, B in ft-L)
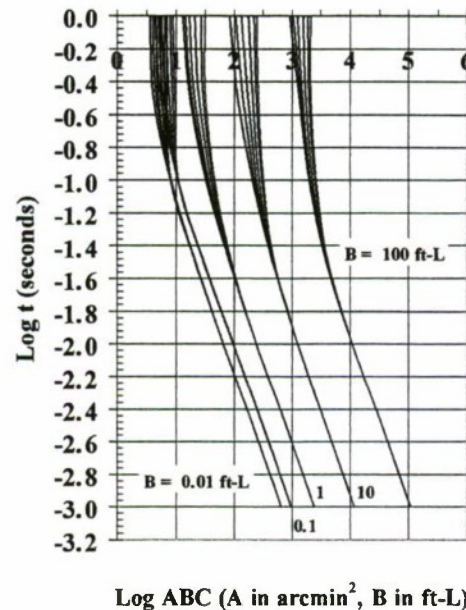
Figure 7. The background luminance family of log t versus log ABC plots after a horizontal co-ordinate translation.

Figure 7 illustrates the merging of several families of hyperbolas (B = 0.01, 0.1, 1, 10 and 100 ft-L) relating to plots of log t versus log ABC produced by a horizontal co-ordinate translation. Those plots with the same background luminance have been horizontally translated to a point of coincidence with the plot corresponding to A = 2825 arcmin^2. The interest in these mergers is that it can now be seen that all plots of log t versus log ABC have an extended section that is an approximate straight-line segment with a slope of -1. For these regions the ABCt law simplifies to

$$\log t = -\log ABC + \log N \qquad (23)$$

where log N is the y-intercept of the approximate straight-line segment. Equation (23) can also be written as

$$ABCt = N \qquad (24)$$

where N is a constant for a given A and B. The parameter N represents the threshold energy required to elicit excitations from the foveal visual system. The duration for which Eq.(24) is valid is called the critical duration and varies with background luminance. It is somewhat affected by target area. It ranges from approximately 30 msec at high background

luminance to approximately 100 msec for low background luminance. It is within these durations that the visual system is most efficient. An equation for N can be obtained from the equations of a straight-line and Eq. (15). Since the line of slope -1 passes through the point $(x^*, -3)$, using $(xN, yN) = (x^*, -3)$ in Eq. (15) yields

$$AN x^{*2} + (DN - 3 AN) x^* + (FN - 3 EN) = 0 \qquad (25)$$

Solving for x^* gives

$$x^* = \frac{(3A' - D') + \sqrt{(D' - 3A)^2 - 4A'(F' - 3E')}}{2A'} \qquad (26)$$

From the equation of a line with slope -1 passing through the points $(x^*, -3)$ and $(0, \log N)$

$$\log N = x^* - 3. \qquad (27)$$

Since AN, DN, EN and FN are all functions of target area and background luminance, log N is also a function of these two parameters. Figure 8 shows the relationship between threshold energy (log N), target area (log A) and background luminance during the critical duration when ABCt = N. An analysis of Figure 8 reveals a number of important correlations. For any target diameter less than one degree, the threshold energy increases with background luminance. However, there is very little difference in the threshold energies at very low background luminance levels (B < 0.01 ft-L). In the Ricco=s law regime ($A < 6$ min^2), threshold energy is independent of target area; threshold energy is constant for a given background luminance. As target area increases beyond the Ricco=s law regime, threshold energy is also increased. For target areas $A > 100$ min^2, plots of log N versus log A appear to be straight, parallel lines with slopes roughly equal to 0.75 (the slopes range from 0.72 to 0.85). This implies

$$\log N = 0.75 \log A + \log K \qquad (28)$$

where log K is the y-intercept and a function of background luminance B. Equation (28) can also be written as

$$N = K A^{0.75} \qquad (29)$$

where A is in arcmin^2, B in ft-L and t in seconds.

In other words, for large targets, the threshold energy is proportional to $A^{0.75}$.
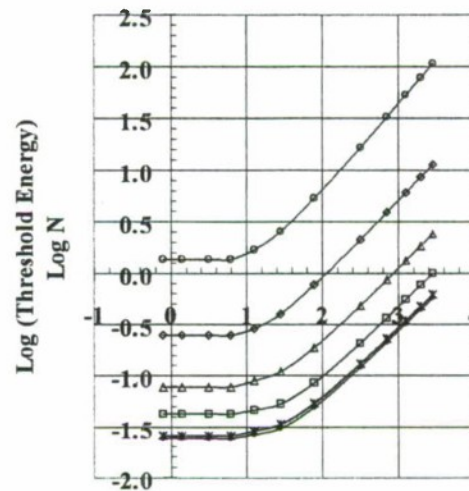
Log A (A in arcmin^2, B in ft-L)

| -•- B = 0.001 | -*- B = 0.01 | -•- B = 0.1 |
| -•- B = 1 | -•- B = 10 | -•- B = 100 |

Figure 8. Temporal summation of light energy. Y-intercept (log N) versus log A.

The following example finds the value of N for A = 78.54 arcmin^2 and B = 10 ft-L: Inserting the values from Eq. (20) into Eq. (26) gives $x^* = 2.88$. From Eq. (27) we get log N = -0.12 and N = 0.76. At t = 0.001 seconds, the BM data set gives C = 0.982; hence, the product ABCt = 0.77.

The above observations relate to Bloch=s law, which can be expressed as follows: the product of threshold intensity and target duration is a constant within the critical duration. During the critical duration we have what is commonly referred to as to as temporal summation of light energy. Total temporal summation of light energy occurs when ABCt = constant. It is important to remember that the value of N in Eq. (24) is a function of both A and B. It is not appropriate to substitute values for A and B without adjusting the value of N. Figure 9 illustrates the effect of using different values for A and B that produce the same product for AB. The plots of log t versus log ABC for the same product AB are not the same.

As target duration increases beyond the critical duration, the slopes of the tangents to the log t versus log ABC curve approach infinity; the visual system becomes increasingly inefficient as noise begins to dominate the system. As the slopes approach infinity, the threshold intensity ABC approaches some minimum value related to the Weber-Fechner law that depends on the target area and the background luminance. After a critical duration, target

249

duration no longer has any influence and brightness discrimination is determined entirely by the incremental luminance of the target. This minimum is easily obtained from the ABCt law by differentiating Eq. (15) with respect to x and allowing the derivative to approach infinity. This procedure yields

$$\log (ABC)_{min} = -EN/AN \tag{30}$$

where AN and EN are given by Eqs. (6), (10), (16) and (17). Figure 10 shows how the minimum intensity varies with target area and background luminance. In the Ricco=s law regime, the minimum threshold intensity is a constant for a given background luminance; smaller values are associated with higher luminance levels. As the target area increases beyond the Ricco area, the minimum threshold intensity also increases; more light energy is required to elicit excitation. It can also be seen from Figure 9 that for a given target area the minimum threshold intensity increases as the background luminance increases. In fact, for target areas greater than approximately 100 arcmin^2, the minimum threshold intensity is directly proportional to the target area. Using the approximation

$$- (B_2 E_2 - 2 C_2 D_2) \cong B_2 \sqrt{E_2^2 - 4 C_2 F_2} \tag{31}$$

in Eq. (16), it can easily be shown that the plots in Figure 10 are hyperbolas with asymptotic slopes of zero and +1.

5. TWO PART FORMALISM OF THE ABCt LAW

Interestingly, the critical duration seems to be related to the response time of the cone cells in the foveal region of the retina. Studies performed on monkeys[7], using measurements of membrane currents from monkey cones generated by brief flashes of light of varying strength, show that their response times correspond closely to the critical duration for humans. An alternate formalism of the ABCt law has also been pursued that closely relates to the critical duration and the response mechanism of the human eye. It considers the threshold presentation time t to consist of the product of two functions, Q and t*, where t* represents the total integration of light energy (ABCt) over time. This is referred to as temporal summation. Total temporal summation occurs when ABCt = constant (during the critical duration). The Q parameter is a measure of the amount that the foveal vision system deviates from total temporal summation of light energy. If t = Q t* then $\log t = \log Q + \log t^*$. Using $y = \log t$, $y_1 = \log Q$ and $y_2 = \log t^*$, Eq. (15) becomes

Log ABC (A in arcmin^2, B in ft-L)

→ A = 12.56, B = 100	⊣ A = 1256, B = 1
⊣ A = 12.56, B = 0.1	⊣ A = 1256, B = 0.001
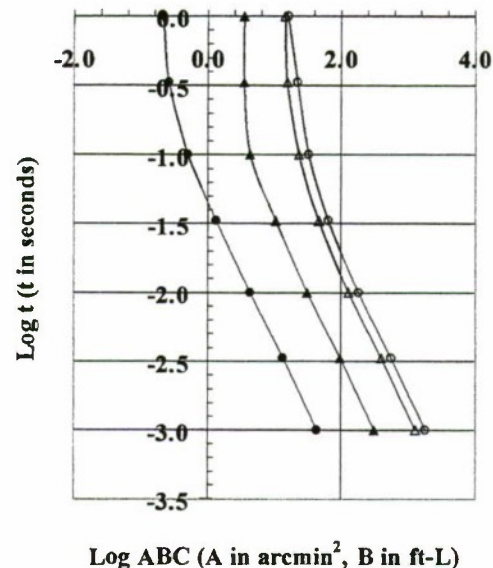
Figure 9. Plots of log ABC versus log t are not unique for the same product AB.

Log A (A in arcmin^2, B in ft-L)

→ B = 0.001	⊣ B = 0.01	→ B = 0.1
⊣ B = 1	⊣ B = 10	→ B = 100
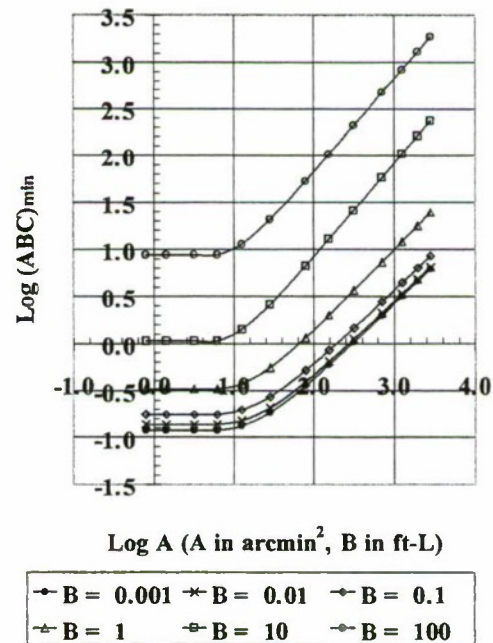
Figure 10. Minimum log ABC versus log A.

$$AN x^2 + AN x (y_1 + y_2) + DN x + EN (y_1 + y_2) + FN = 0 \tag{32}$$

If y_2 is the y-coordinate of a straight line through the point (x*, -3) with slope = -1, then

$$y_2 = -x + x^* - 3 \qquad (33)$$

and Eq. (32) can be written as

$$AN \; x \; y_1 + EN \; y_1 - (3 \; AN - DN - AN \; x^* + EN) \; x$$
$$- (3 \; EN - EN \; x^* - FN) = 0$$
$$\qquad (34)$$

Since closed-form equations have been derived for the y-intercepts ($y_0 = \log N$), Eq. (34) represents a closed-form solution for the multiplicative function Q. Also, since the coefficients of second degree vanish, Eq. (34) is the equation of a hyperbola with one asymptote vertical and the other horizontal. Figure 11 illustrates the two-part formalism of the ABCt law. Evidently, the process of detecting a BM target, using foveal vision, can be equated to a linear process associated with total temporal summation of light energy modified by a quadratic multiplicative function.

6. RESULTS

The validity of the ABCt law can be quickly determined by comparing BM threshold contrasts with those calculated using the ABCt law. Of the 420 unique BM threshold contrasts, 21 threshold contrasts obtained from the ABCt law produce percent differences greater than 10 percent. These contrasts are identified in Table 1 along with the conditions that produced them. Seven of the Table 1 entries differ from the BM threshold contrasts by less than 0.10. Also, the largest percent difference is less than 17 % with the majority of the errors in the 11-13 % range. Hence, only three percent of the unique BM threshold contrasts are reproduced with any appreciable error.

$$A = 78.5 \; \text{arcmin}^2$$
$$B = 10 \; \text{ft-L}$$
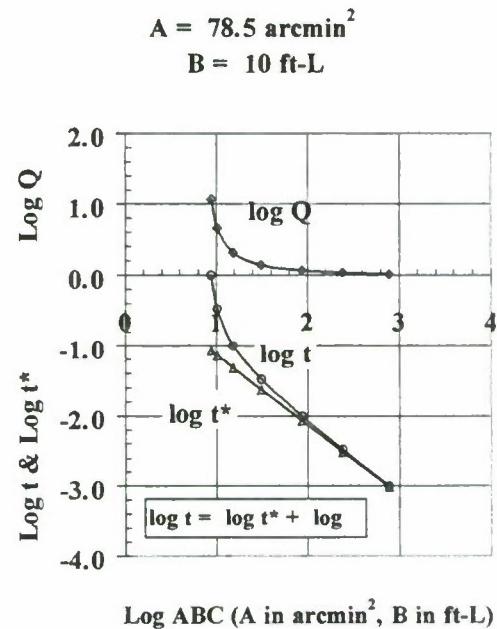
Figure 11. Two part formalism of the ABCt law.

Table 1. Conditions producing a percent difference greater than 10% between BM threshold contrasts and those calculated from the ABCt law.

#	A (min^2)	BM C	Model C	C % Diff	Δ C	t (sec)	B (ft-L)
1	314.12	0.020	0.022	10.90	0.00	1/30	100
2	706.46	0.017	0.019	10.65	0.00	1/30	100
3	28.25	0.320	0.285	-10.94	0.03	1/300	100
4	28.25	1.035	0.900	-13.03	0.13	1/1000	100
5	78.54	0.764	0.683	-10.54	0.08	1/1000	100
6	153.85	0.082	0.092	11.35	-0.01	1/100	10
7	314.12	0.068	0.075	10.39	-0.01	1/100	10
8	706.46	0.113	0.132	16.57	-0.02	1/100	1
9	6.15	1.941	2.186	12.62	-0.24	1/30	0.1
10	6.15	5.433	4.560	-16.05	0.87	1/10	0.01
11	6.15	13.740	12.207	-11.16	1.53	1/30	0.01
12	6.15	134.276	115.931	-13.66	18.35	1/300	0.01
13	314.12	13.274	11.700	-11.86	1.57	1/300	0.01
14	706.46	10.257	9.071	-11.55	1.19	1/300	0.01
15	6.15	434.510	382.474	-11.98	52.04	1/1000	0.01
16	314.12	42.954	38.637	-10.05	4.32	1/1000	0.01
17	6.15	24.099	21.337	-11.46	2.76	1	0.001
18	6.15	52.723	47.186	-10.50	5.54	1/10	0.001
19	12.56	205.589	232.681	13.18	-27.09	1/100	0.001
20	28.25	114.815	130.145	13.35	-15.33	1/100	0.001
21	2825.46	20.370	22.685	11.36	-2.31	1/100	0.001

REFERENCES

1. R.M. Matchko and G.R. Gerhart, AParametric analysis of the Blackwell-McCready data,@ Opt. Eng. 37 (7) 1937-1944 (1998).
2. R.M. Matchko and G.R. Gerhart, Errata, Opt. Eng. 38 (3) 564 (1999).
3. R.M. Matchko and G.R. Gerhart, ATarget detection for low ambient backgrounds,@ Opt. Eng. 37 (7) 1945-1950 (1998).
4. H.R. Blackwell, AContrast thresholds of the human eye,@ J. Opt. Soc. Am. 36, 624-643 (1946).
5. H.R. Blackwell and D.W. McCready, A Foveal contrast thresholds for various durations of single pulses,@ Report # 2455-13-F, U. Michigan Visual Research Laboratories (1958).
6. R.J. Farrell and J.M. Booth, *Design Handbook for Imagery Interpretation Equipment*, Boeing Aerospace., Seattle, Sec. 3.2, p. 8 (1984)
7. J.L. Schnapf and D.A. Baylor, AHow the photoreceptor cells respond to light,@ Sc. Am., 256 (4): 40-47 (1987).

A Model to Evaluate the Decision-Making Process in a Target Acquisition Task

Masha Maltz

Ben-Gurion University of the Negev

Department of Industrial Engineering and Management

ABSTRACT

We consider in this paper observer behavior when images from different sensors both with and without targets are presented for a target acquisition task. False alarm rates are shown to increase with the change from a regular to a zoom sensor. In addition, the presence of an obvious target inhibits false detections. Implications for target acquisition models are discussed.

I. INTRODUCTION

In most applications of target detection electro-optical systems, the man-in-the-loop plays a crucial part in the performance of the system. If one were dealing with an automatic target acquisition system, a great deal of effort would justifiably be invested in calculating the threshold for determining a target for optimum system performance. When the system is human-operated, understanding the decision-making process of the human observer is of major importance to the successful performance of the system.

The question arises: how can we model the human observer? The theoretical foundation for modeling the human decision-making process in a detection task has been advanced with the development of Signal Detection Theory (SDT). SDT explains the behavior and the detection ability of the observer in terms of the tradeoffs between false alarms and missed signals that generally characterize binary decisions. Within the framework of SDT, the observer's hits, misses, false alarms, and correct rejections are combined to give an overall view of the observer's decision strategy as well as performance ability.

While the model exists, the determinants of specific observer behavior are still open to question. In particular, how does the observer modify his behavior when there are changes in the number of false alarm opportunities? On this question, there has recently been a debate in the literature. Grossman et al (1) argued that the observer attempts to maintain a constant false alarm rate (CFAR), so that when the number of false alarm opportunities increases, the observer will correspondingly raise his/her decision criterion and reject more of them. Doll and Schmieder (2), on the other hand, believed that the human observer does not necessarily reject the possibility of increased false alarms when there are more opportunities.

Both of the above studies based their findings on relating the amount of clutter in an image to the number of false alarms measured. In this paper, we wish to expand this concept to determine whether observer performance remains consistent under the conditions presented by two additional variables: 1) the case where two different sensors are used to present the stimuli to the observer, and 2) the case where some of the images have obvious targets while others have no targets.

In Section II, we briefly describe signal detection theory. We describe our experiment and its results in Section III and IV. A discussion of the implementation of these results into a general model of target acquisition is presented in Section V.
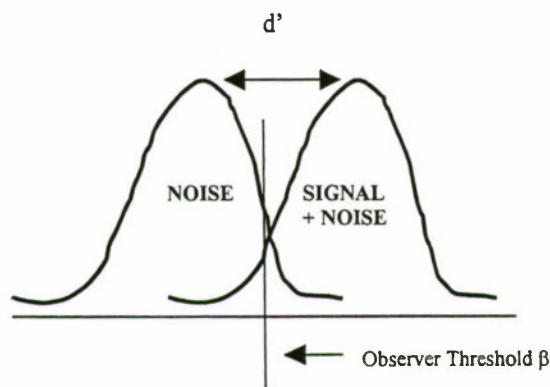
II. SIGNAL DETECTION THEORY

Signal Detection Theory (SDT) can be briefly summarized for our purposes as follows (see reference 3 for a detailed description):

An observer is presented with a serious of opportunities to determine the presence of a target. If a target is present and it is detected, a "hit" is registered; if it is not detected, then a "miss" occurs. Similarly, if only noise is present and a target is mistakenly detected, a "false alarm" has occurred. If the noise is correctly recognized as noise, it is defined as a "correct rejection."

When noise is presented to the observer, there exists a probability distribution associated with the degree that the noise resembles a target. For the noise curve in Fig. 1, the 'x' axis is the degree of similarity between the noise and a real target and the 'y' axis is the probability value that the noise will have this value. For the case when a target with accompanying noise is presented, its distribution of values is related to how similar it looks to a clear target.

When the two distributions overlap, the observer, who places a threshold and decides that above such a value he will declare a target, can never achieve a perfect result (0% misses and 0% false alarms). The observer's ability to discriminate between noise and target with noise is limited by the distance between the means of the two distributions, defined as the variable d'. Shifts of the threshold, defined as the variable β, will result in changes in the tradeoff between misses and false alarms. Observers wishing to avoid missed targets will adjust their threshold downward and will, as a result, accept more false alarms. The placement of the threshold is governed by the conscious or unconscious decision of the observer. How the position of the threshold changes with different scenarios is the subject of this paper.

Fig. 1: Probability distributions for noise and signal (target) plus noise as a function of similarity to an ideal (noise-free) target.

III. EXPERIMENTAL PROCEDURE

Five infrared pictures, with one target each, were prepared. Four of the images' targets were easily recognizable, and one image's target was, in our opinion, difficult to find. Two variations of these images were prepared. A quarter of each of the original images was cut out and blown up to fill the original image size. The target was inserted into the resultant image in such a way as to avoid artificial cues to the presence of the target. For the purpose of this paper, we will consider such images to come from a "different" type of sensor, i.e. we will refer to the "regular" sensor and the "zoom" sensor. These ten images (five regular and five zoom) were then duplicated with their targets removed.

The twenty images were displayed on 17-inch computer screens to fifteen observers in a laboratory setting. The observers were not informed of how many targets were present in any image. The order of image presentation was counterbalanced between subjects, and ordered so that the duplication of the images would have little or no effect. A post-experiment survey confirmed that the recurrence of similar images was not apparent to the subjects.

Observers were told to move a mouse-controlled cursor to the target positions that they found and to "shoot" at the targets by clicking on the mouse button. The location and time of the clicks were recorded for our analysis.

IV. RESULTS

Table 1 gives the probability of a hit for the target in each of the images with targets and the average number of false alarms per observer for each image. The probability of hit was fairly constant for the two sensors. However, the number of false alarms per image increased from 0.23 to 0.41.

To check the significance of these results, we ran the statistical test, analysis of variance (ANOVA) on the false alarm rate, testing for the variables sensor type and the target presence (or absence). We found that both sensor type and target presence were significant main factors affecting false alarm rate. The interaction between sensor type and target presence was not significant. Figure 2 shows the dependence of the number of false detections

per observer for images on the regular and the zoom sensors, and figure 3 shows the effect of target presence.

Table 1: Probability of detection and false alarm rates for images with targets from the two sensors.

	Sensor 1 (regular		Sensor 2 (zoom)	
	Hits	False alarms	Hits	False alarm
Image 1	0.8	0.3	1.0	0.4
Image 2	0.9	0.0	1.0	0.3
Image 3	0.5	0.5	0.3	0.7
Image 4	1.0	0.1	1.0	0.6
Image 5	1.0	0.1	1.0	0.1
	0.84	0.23	0.87	0.41
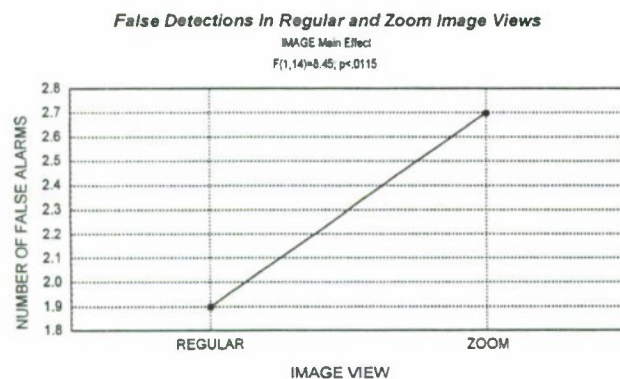
Fig. 2: Average false alarm rate per observer (summed over ten images) for images from the regular sensor and the zoom sensor.
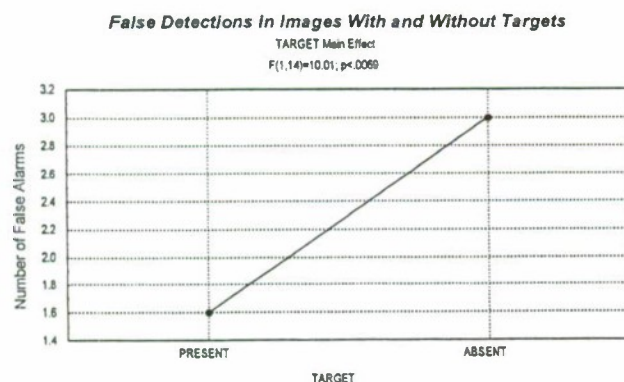
Fig. 3: Average false alarm rate per observer (summed over ten images) for images with targets present and those without targets.

V. DISCUSSION

Sensors: Our results when comparing false alarms rates under the regular and the zoom sensors, showed that observers did not strive to maintain a constant false alarm rate. The zoom sensor had a significantly higher false alarm rate than the regular sensor when comparing all the images. This would indicate that the CFAR model of observer behavior is limited and does not extend to major changes between sensors.

It is interesting to note that the change in the false alarm rate cannot automatically be equated with a higher false alarm probability. To determine the false alarm probability, one must calculate how many opportunities existed in the images. We chose two different methods for calculating the number of false alarm opportunities to illustrate how the interpretation of the results can differ drastically with the method used. One of the methods was suggested by Grossman et al [1], whereby there are a fixed number of false alarm opportunities in each image based on dividing the area of the image into glimpse-sized sections. Under their method, a higher false alarm rate would indicate a lower value for d', or decreased observer ability to discern between targets and noise.

A second method that we used was to include as opportunities the actual areas chosen by the subjects as false alarms within the images. This method is similar to one used by Doll et al [2] in that observers' actual decisions were used post-facto to establish the pre-experiment circumstances. In our calculation, we defined as false alarm opportunities those areas designated by two or more different observers as containing a target. In this way, we avoided basing our calculations on irrational decisions by any individual observers. Under this method, a higher false alarm rate could be related to more false alarm opportunities and not necessarily to decreased detection ability or to more difficult images.

We tried both explanations in the analysis of our data. Table 2 contains the resultant d' calculations based on the two methods. In the first, we considered a false alarm opportunity to be any area that any two (out of the fifteen) observers incorrectly considered to be a target. In the second, we assumed a constant average of three false alarms per image. As shown in the table, d' increases with the first method; there are more false alarms not because the opportunities are more target-like, but because there are more plentiful. Alternatively, with the

255

second method where we assume that the number of false alarm opportunities is fixed, d' decreases while the number of false alarms increases.

Table 2: Value of d' for images from the two sensors based on two alternative methods of calculating the number of false alarm opportunities, as described in the text.

False alarm opportunities and d'		
	d'(a)	d'(b)
All Images	2.11	2.36
sensor 1 (Regular)	2.12	2.41
sensor 2 (Zoom)	2.17	2.26
a: false alarm opportunities > 1 observer b: given a fixed # of opportunities (3/image)		

Target Presence: As shown earlier in figure 3, we show the effect on observer false alarm rate by target presence. In this effect, again, the CFAR model is rejected. The presence of an obvious target affected our observers' tendency to detect false alarms. The average number of false alarms per observer for images without targets was 3.0 compared to 1.6 for images with targets. Interestingly, the figures in table 1 point out that for the image with the difficult-to-discern target (Image 3) there was an increased number of false alarms. Thus, while, obvious targets seem to inhibit false alarm detections within the image, difficult targets seem to bolster false alarm rates.

CONCLUSIONS

This study was performed to study the differences in observer performance when confronted with different sensors and different conditions. We found that CFAR was not maintained when going from one sensor to the next, and that there was a tendency on the part of observers to lower the number of false alarms when an obvious target was present and detected.

The increase of the number of false alarms with a change in sensor has at least two possible explanations: 1) the number of false alarm opportunities increased or 2) the similarity of false alarms to real targets increased with the

usage of the zoom sensor. Differentiating between these options is an interesting direction for future research.

A good model of human target acquisition is necessary to be able to predict performance with the electro-optical systems discussed here. We hope that this paper represents a small contribution in this direction.

VI. REFERENCES

[1] Grossman, S., Hadar, Y., Rehavi, A., and Rotman, S.R., 1995, "Target Acquisition and False Alarms in Clutter," *Optical Engineering*, Vol. 34, No. 8, pp. 2487-2495.

[2] Doll, T.J. and Schmieder, D.E., 1993, "Observer False Alarm Effects on Detection in Clutter," *Optical Engineering*, Vol. 32, No. 7, pp. 1675-1684.

[3] Green, D.M. and Swets, J.A., 1966, *Signal Detection Theory and Psychophysics*, John Wiley & Sons: New York.

Image-Based Model for Visual Search and Target Acquisition

Magnús Snorrason[a], Harald Ruda[b], and James Hoffman[c]

[a,b]Charles River Analytics Inc., Cambridge, MA 02138, USA

[c]Department of Psychology, University of Delaware, Newark, DE 19716

ABSTRACT

The overall objective of this project is to develop a computational model for predicting probability of detection during search for hard-to-see targets. This model is image based: it uses imagery for input, rather than estimated parameter values characterizing critical factors such as clutter and target detectability. Consequently, it generates probability-of-detection values that are functions of image content, rather than functions of subjectively estimated parameters. The input domain is infrared or visible-light imagery of distant vehicle targets in cluttered scenes. Such hard-to-see targets are generally only detected once they have been fixated. Hence, our modeling approach focuses primarily on factors influencing the choice of fixation points during visual search. A saliency map is constructed from bottom-up image features, such as local contrast. To account for top-down cognitive effects—such as bias towards the horizon—a separate cognitive bias map is generated. The combination of these two maps provides a Fixation Probability Map (FPM). Given the FPM, a sequence of fixation points is generated in a way that accounts for imperfect memory of past fixation locations. Results are presented comparing model-generated FPM's with eye-tracker data collected from observers in visual search experiments.

INTRODUCTION

The goal of this project is to construct a model that elucidates the relationship between bottom-up (stimulus driven) and top-down (cognitive) effects on visual search for distant hard-to-detect targets, with the ultimate goal of predicting the average search performance of a trained operator on a given scene.

The task of visual search—locating objects of a known category based on their visual appearance—is among the many functions of the human visual system that are not well understood. Given the multitude of stimulus variables that affect visual search, modeling this task is not possible without further constraints. Past approaches have generally been constrained by fitting psychophysical data from very simple scenes: uniform background with an array of many similar distractor items and one or a few target items (for review see Wolfe, 1998). Unfortunately, such models do not necessarily scale with increased scene complexity and are not applicable to photo-realistic imagery. Models specifically designed to handle photo-realistic imagery, such as the TARDEC vision model (Gerhart, et al, 1995), Georgia Tech vision model (Doll, et al, 1997), and recent work at CalTech (Itti, Koch, and Niebur, 1998) have not included cognitive effects. There is great need for models capable of handling imagery of realistic complexity, yet designed with modeling constraints that enable both bottom-up and top-down factors of visual search to be better understood.

We chose to constrain the problem by focusing on hard-to-see targets, targets that can not be detected without foveation (where we define "detection" as a true positive response in a visual search experiment, i.e. designating a real target as "target"). This is a reasonable constraint because detecting a hard-to-see target requires high visual acuity, and visual acuity is a function of retinal eccentricity. Acuity is highest in fovea with rapid fall-off towards the periphery. Acuity is reduced by 50% just 2.5° away from the point of fixation and it is reduced by 75% at 6° eccentricity.

Acuity falloff is nature's way of solving a resource allocation problem: the fovea is far more useful than the periphery, but high acuity requires resource intensive high-resolution sampling and processing of the image. In fact, it is so resource intensive that even though foveal field-of-view covers only a region the size of the tip of your thumb at an arms length, a large portion of the human brain is dedicated to foveal processing.

With such a small fovea, most of the visual scene never gets foveated. This led us to the conclusion that *the choice of fixation points in a given scene is the most important determinant of whether hard-to-see targets get detected.* Hence, our modeling approach focused primarily on the different factors that affect the choice of fixation points and secondarily on the actual foveal processing.

MODEL

We have developed a model with four main components: *Peripheral Processing, Fixation Selection, Foveal Processing,* and *Classification.* The block diagram in Figure 1 below shows the relationships between those components. The following table summarizes the functionality of each component. The whole retinal image is continuously sampled at low resolution by the *Peripheral Processing* component while at the same time one small region centered on the fixation point is processed in full resolution by the *Foveal Processing* component. The *Fixation Selection* component processes a *Fixation Probability Map* from the peripheral component to generate coordinates specifying the next fixation point. The output from the foveal component is in the form of feature vectors that get classified by the *Classification* component, such that each fixation point gets labeled "target" or "non-target".
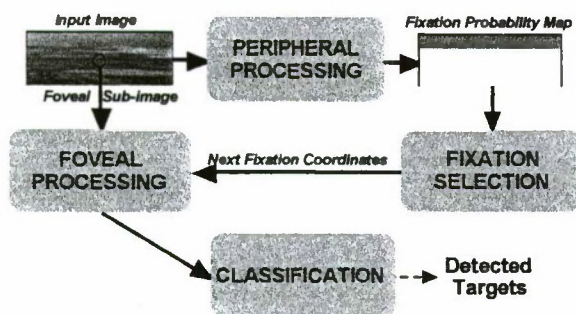

Figure 1: Block diagram for search and detection model

Table 1: Summary of main component functionality

Module	Description
Peripheral Processing	Computes Fixation Probability Map using bandpass filtering followed by a nonlinear combination of local contrast measures and an optional bias towards the horizon
Fixation Selection	Computes fixation coordinates based on highest peaks within Fixation Probability Maps, accounting for imperfect memory of past fixations
Foveal Processing	Computes feature-data for the neighborhood of each fixation point, based on expected target size and shape
Classification	Based on the features, classifies each fixation as a point where human observers would either designate a target or a non-target

The following subsections describe the first two components in more detail, the rest of the model will be reported elsewhere.

Peripheral Processing

The *Peripheral Processing* component receives input from the peripheral regions of the retinal image. Since that is most of the image, the simplifying assumption was made to use the whole image as input to Peripheral Processing. The output of this component is a Fixation Probability Map (FPM) of the same size as the input image, where the value of each pixel indicates the predicted probability for a fixation at that location in the input image. The following modules operate sequentially within this component:

- Bandpass Filtering of the whole image models retinal lateral inhibition and low-resolution sampling.
- Peripheral Feature Extractor computes Peripheral Feature Maps of the same dimensions as the input image, one map per feature.
- Threshold & Sum operator combines all feature maps into one Saliency Map.
- Horizon Detector generates a Horizon Bias Map if the horizon is visible in the image.
- Sum & Normalize operator combines the Saliency Map and Horizon Bias Map into a Fixation Probability Map.

Bandpass Filtering

A modeling approach was needed for the peripheral low-resolution sampling of the input image, that also captures the well known lateral inhibition found in many processing layers of the retina and early cortical regions. We chose a difference-of-Gaussians (DOG) convolution operator with parameters (center and surround Gaussian widths and heights) chosen to emulate known physiological and anatomical data (Spillman & Werner, 1990, p.29). The resulting DOG consists of two concentric Gaussians. Modeling lateral *inhibition*, the surround Gaussian is *subtracted* from the center one. The width of the center Gaussian was chosen to represent retinal sampling at around 6° eccentricity, with $\sigma^2 = 2$. The surround Gaussian was chosen to be 5 times wider and 1/5 as high. Convolving this DOG with an image blurs the image slightly (due to the center Gaussian), but it also high-pass filters it (due to the inhibitory surround), with a net effect of bandpass filtering as shown in Figure 2 below. This increases the salience of very low contrast targets and helps explain how such targets attract fixation.
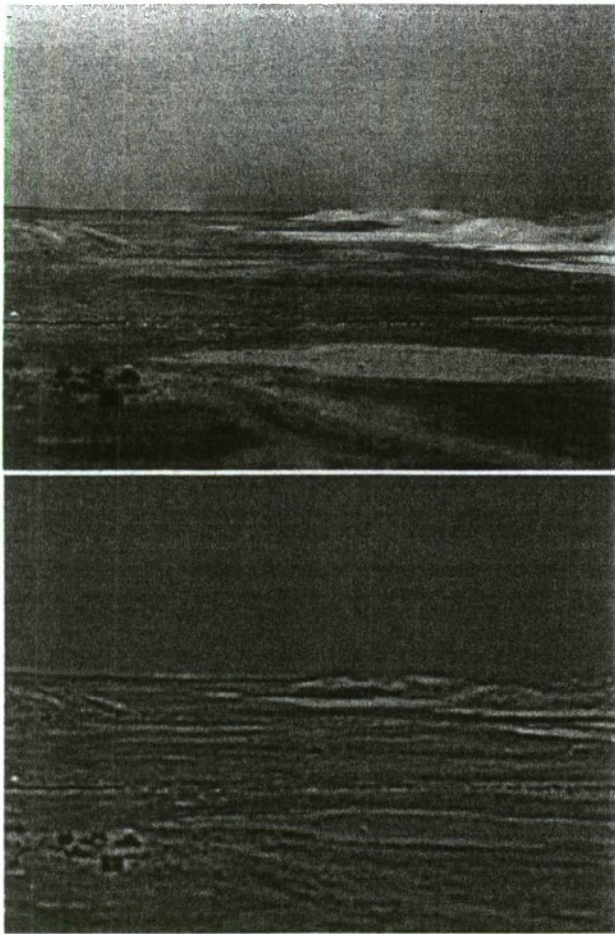
Figure 2: Original input image (above) and bandpass filtered version (below)

Peripheral Feature Extraction

We assume that the probability of attracting attention to any given image region depends to a large degree on the feature *saliency* of that region with respect to the associated background. The Peripheral Feature Extractor measures feature saliency by computing Peripheral Feature Maps (PFM) from the filtered input image.

The concept of a feature map is attractive for a number of reasons. Feature maps are biologically plausible: they are topographic representations of the retinal image, just like the maps that have already been identified in visual cortex of primates and cats. In fact, there is significant evidence that topographic maps are the primary representation for *all* sensory information in the brain. Feature maps are fully parallel, both across spatial locations within each map and across all maps. Such parallelism is needed to explain the very rapid effects of visual "pop out", whereby one can detect a target object among distractors in as little as 50ms (after response delays have been

factored out). Finally, from a modeling point-of-view, feature maps are easy to work with. They retain all spatial relationships (unlike frequency domain representations) so they can always be viewed as images.

Using feature maps begs the question: what are the features? After detailed analysis of past feature integration approaches (see Wolfe, 1998 for a review), we decided that *local contrast* was the feature with the best potential in our domain, i.e. visual search for small hard-to-see targets in gray-scale still images. We implemented and evaluated a number of local, unoriented contrast measures. Through analysis and empirical testing, we found the following four contrast measures most successful:

- Absolute-difference-of-medians
- Dispersion
- Difference-of-dispersions
- Doyle

The **absolute-difference-of-medians** contrast saliency measure at a point is defined as the absolute difference between the median brightness of a rectangular window A_1 centered around the point P and the median brightness of a larger rectangular window A_2, as illustrated in Figure 3 below, and computed via equation (1):

$$C_P = |\ Med(A_1) - Med(A_2)\ | \qquad (1)$$

where C_P is the contrast metric at point P and $Med(A)$ is the median brightness value inside area A. The median value is chosen over other low pass filters due to its preservation of edges. Other filters, such as the Gaussian filter or other linear smoothing filters, were considered for this function, but were found less desirable since they result in the presence of "shadows" near the edges. An additional advantage to using the median is that for standard byte-per-pixel images, the median value can be computed inexpensively using a histogram technique.
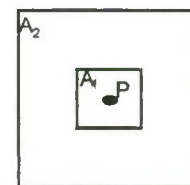

Figure 3: Two-window contrast measure kernel

This contrast measure depends on the choice of the two free parameters areas A_1 and A_2. An object which matches the size of A_1 will give maximum contrast response. Since the size of objects of interest in the image

259

is not known a priori to the operator, we implemented the contrast filter with multiple different sizes for A_1 computed in parallel. The size of A_2 was kept constant since it was significantly larger than even the largest A_1. This is effectively equivalent to having multiple contrast filters operate in parallel. The outputs of the different contrast filters are then averaged to generate an overall contrast measure. The reason for taking the absolute value of the difference was to make the measure equally sensitive to either contrast polarity, i.e., light object on a dark background or dark object on a light background.

Dispersion is the sum of absolute values of differences between the pixel intensity at point P and at all other points inside window A centered on P. This is an L1 measure of the local variance in intensity values around a pixel, in the same way as σ^2 is the L2 measure of local variance. (In fact the dispersion and difference-of-dispersions features can be replaced by variance and difference-of-variance features without major adverse effects). This measure of local contrast has one free parameter: the window size A. Unlike the difference-of-medians, this method is sensitive to intensity variations *inside* the window, hence only one window is used. For small window sizes, the dispersion value is a measure of abrupt variation in intensity levels, such as caused by edges running at any orientation through the window.

We also implemented a **difference-of-dispersions** contrast saliency measure using the same two window approach shown in Figure 3 above. The two window sizes are combined by subtracting the dispersion value for the larger window from the one for the smaller window and half-wave-rectifying the result. This is equivalent to saying "show me the locations where dispersion on a small scale is *greater* than dispersion on a large scale." This is a selective measure of contrast, but one that correlates well with some highly salient image regions.

An efficient way of computing dispersion used for both dispersion contrast measures is based on the histogram of pixel intensities computed over window A. The dispersion at point P is then given by:

$$D_p = \sum_{k=1}^{n} \left| z - \mu \right| p(k) \tag{2}$$

where N is the number of histogram bins (normally chosen to be 256), z is the bin center value, p(k) is normalized histogram frequency at bin k, and μ is the mean intensity value inside the window.

The **Doyle** contrast measure is a measure of first-order statistics that combines mean and standard deviation as follows:

$$[(\mu_1 - \mu_2)^2 + (\sigma_1 - \sigma_2)^2]^{1/2} \tag{3}$$

where μ_i and σ_i are mean and standard deviation computed over area A_i, referring once again to windows A_1 and A_2 in Figure 3 above. This measure—unlike the difference-of-dispersions—is equally responsive to either contrast polarity. It responds equally well to a small high contrast region in a larger low contrast area and a small low contrast region in a larger high contrast area. This insensitivity to contrast polarity is similar to the absolute-difference-of-medians, but that method is only sensitive to variations larger than the smaller window, the Doyle method is also sensitive to intensity variations inside the smaller window due to the standard deviation terms.

Threshold & Sum Operator

The Threshold & Sum operator thresholds each Peripheral Feature Map (PFM) with an empirically determined threshold value and then sums the results into one Saliency Map (SM), as shown in Figure 4 below.

The notion behind thresholding each PFM is well established in modeling human perception: small differences in local features such as color or gray-level *do not* contribute to stimulus driven saliency (Wolfe, 1994). A small difference in gray-levels translates to small values of contrast, hence a threshold is used to null out all regions of each feature map with "small" values.

The threshold also serves another purpose: all feature map values greater than the threshold are replaced by a uniform value. This is to model *categorical* perception, the highly non-linear response found in most forms of human perception (Wolfe, 1994). A thresholded feature map can therefore be interpreted as a map that specifies for each pixel location whether that feature is peripherally perceived as "present" or "absent."

Combining the thresholded feature maps is done by point-wise summation, similar to the original feature integration theory models of Treisman & Gelade, 1980. All feature values that survive thresholding in the individual PFM end up contributing to the Saliency Map. If a particular image location activated multiple PFM strongly, then that location will have a high value in the SM. In other words, that location is highly salient and very likely to attract attention. The SM in Figure 4 below is the result of thresholding, summing, and smoothing the four PFM generated by applying the four feature types described above to the image shown in Figure 2 above. Note how

the highly salient objects in the lower left corner are lightest in this saliency map.
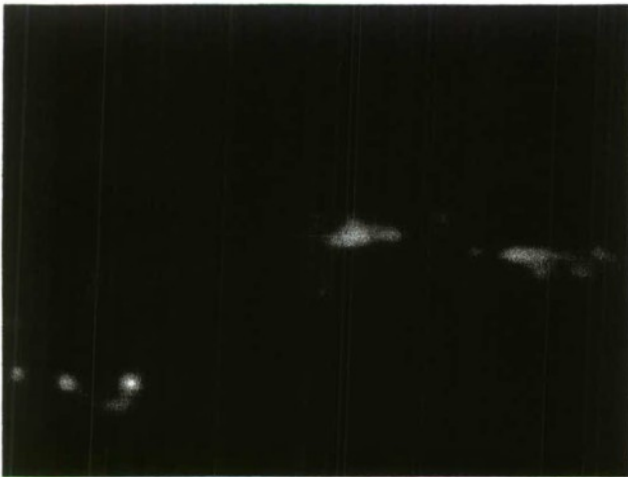

Figure 4: Saliency Map

Horizon Detector

The Saliency Map is the basic component in the Fixation Probability Map (FPM). Without the optional Horizon Detector, the FPM is simply a unity-volume normalized version of the SM, such that the sum of all values in the FPM equals 1.0. However, we recommend including the Horizon Detector for imagery where the horizon is visible.

To summarize our reasoning for including the Horizon Detector, when evaluating FPM based solely on Saliency Maps, we noticed that they consistently underestimated fixation probabilities near the horizon. The eye-tracker data showed that fixations were often densest in a horizontal band bounded by the horizon from above and tapering down towards the bottom of the image. This observed distribution of fixations coincides with the expected *a priori* probability distribution of targets for relatively flat outdoor scenes, such as these scenes. There is therefore reason to believe that this is a "real" phenomenon and not just an artifact in this particular data-set.

Of course, the model's threshold values for the individual Peripheral Feature Maps could be adjusted such that this region contributed more to the SM, but only at the cost of other regions (near the bottom of the image and even above the horizon) becoming significantly over-represented in the SM. This is an inherent problem with global thresholds, they can not selectively affect some regions more than others. Our solution was to develop a Horizon Detector that generates a Horizon Bias Map as shown in Figure 5 below.
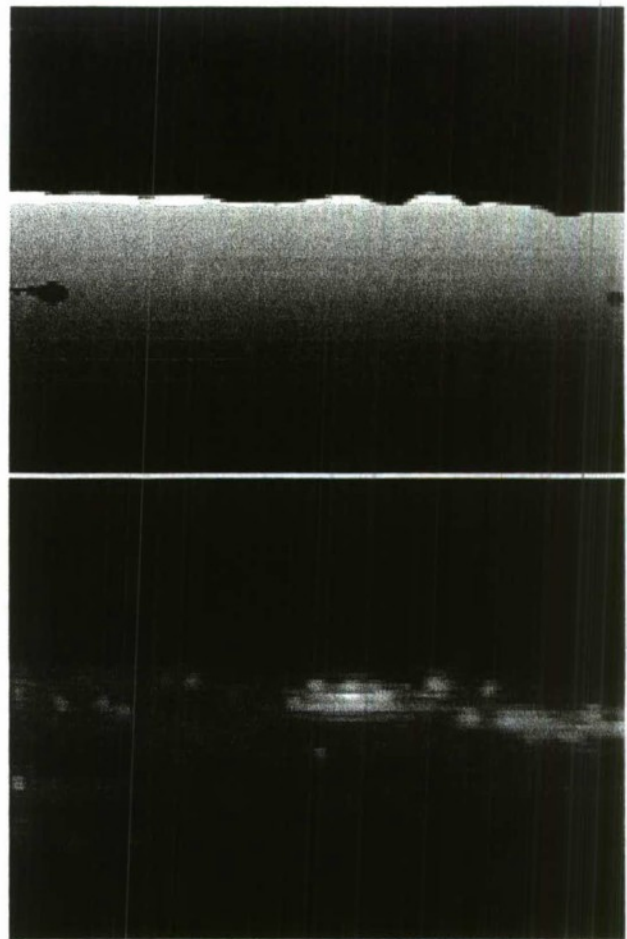
Figure 5: Horizon Gating Map (above) and Horizon Bias Map (below)

Also shown in Figure 5 above is an intermediate step in the making of the Horizon Bias Map. The Horizon Gating Map captures the distribution of fixations described above. It is automatically generated by the following algorithm:

1) Combine PFM into a single image (point-wise max across all PFM)
2) Threshold to binary (black < 10% max)
3) Remove noise (15x15 median filter)
4) Merge small regions until 2 contiguous regions remain (black above horizon, white below)
5) Fill in white region with a linear gradient (white at the horizon to black 10% from the bottom)

This gating map can not be used directly to gate the Saliency Map because that would drastically decrease the effect of highly salient regions close to the bottom of the image (such as the ones in the lower left corner of this image). It can also not be directly added to the Saliency Map because that would imply high fixation probabilities near the horizon even when there are no salient objects there. Consequently, we chose to use it as a gating map for the Peripheral Feature Maps, producing the Horizon

Bias Map shown above. The final result of the Peripheral Processing component—the Fixation Probability Map—is then generated by adding this Horizon Bias Map to the Saliency Map from before and unity-volume normalizing the result (such that the sum of all probabilities is 1.0). Note that the region just below the horizon is now represented as likely to receive numerous fixations, as well as the highly salient object in the bottom right corner.

Fixation Selection

The Fixation Probability Map (FPM) produced by Peripheral Processing is the input for the next major model component, *Fixation Selection*. This component's purpose is to convert the 2-D probability map into an ordered list of fixation coordinates, with the location of highest probability listed 1st, next highest probability 2nd, etc. An obvious solution would be to simply order the peaks in the FPM by height, but that approach would not account for refixations. We have observed in eye-tracker data a small but significant number of refixations for all observers in all but the very simplest search imagery.

Hence, the following process was developed to generate individual fixations from a FPM, accounting for the phenomena of refixation. It involves a feedback loop of the following three modules:

- A Peak Detector chooses the current maximum value in the FPM as coordinates for the next fixation
- A Fixation Memory module subtracts a Gaussian hole from its Fixation Memory Map, while also making all previous such holes in the map slightly shallower
- The holes in the fixation memory map are added to the FPM, preventing the current location and other recently fixated locations from being chosen as the peak value on the next loop iteration

A feedback loop is needed because the choice of future fixations depends on past fixation locations. In particular, the probability of refixating the same location is modeled here with exponential memory decay. The Fixation Memory generates a Fixation Memory Map that represents both the *locations* and *recency* of past fixations, as shown in the following figure. A Gaussian "bump" of fixation probability centered on every recently generated fixation point is *subtracted* from the current Fixation Probability Map, to eliminate the probability of just picking the same fixation point over and over again. To allow for refixations, the memory has a decay term, such that if the peak value of the Gaussian for the current fixation is 1.0 then the peak value for the last fixation is $1.0 - \alpha$, for the next-to-last fixation $1.0 - \alpha^2$, etc., where α

is a free parameter ranging between 0.0 and 1.0. A value of 0.0 for α indicates the assumption of perfect memory (never fixate the same location more than once); $\alpha = 1.0$ indicates an memory-less system (always go back to the last fixation). A value of 0.2 for α was used in simulations shown below.
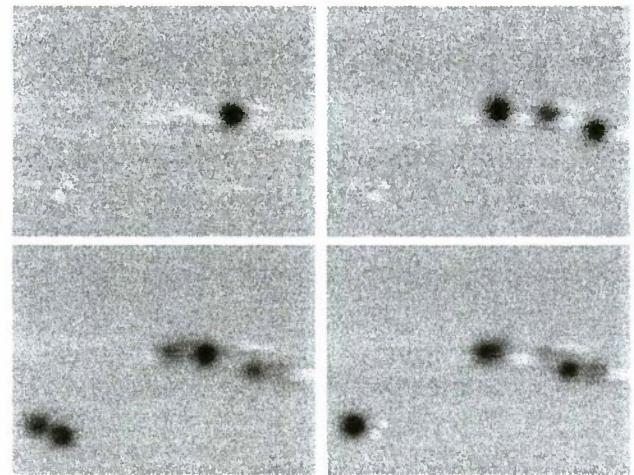
Figure 6: Instantaneous Fixation Probability Maps after 1, 3, 6, and 8 fixations (clockwise from upper left)

Figure 6 above shows the evolution of the Fixation Probability Map as new fixations are generated. Note that the range of values is no longer [0.0, 1.0] because the Gaussian "holes" can be as deep as –3.0. Since gray-scale is still linearly allocated for viewing purposes, black now represents the bottom of the deepest hole in each image and hence the value 0.0 is now represented with a light-gray. (White still represents 1.0.)

The sequence of the past few fixations can be judged from the darkness of the Gaussians in Figure 6 above. For example, after three fixations (upper right panel) it is obvious that the left one is the 3rd (current) fixation, the right one was the 2nd fixation, and the middle one was the 1st fixation. The 6th fixation (lower right panel) is the first fixation to stray from the region of distant hills and the 8th fixation (lower left panel) returns to that region because the Gaussians there have decayed sufficiently. As the Fixation Selection loop is run for more fixations, eventually all regions of significant probability are fixated, with the highest probability regions receiving more fixations. It is also possible to add low levels of Gaussian distributed noise to each pixel value just prior to Fixation Selection, such that a few relatively low probability locations tend to also be picked for fixation.

RESULTS

Evaluating the model's ability to predict fixation locations required access to images and associated eye location data collected with an eye-tracker from observers searching for

targets in those images. We obtained a set of 20 images from the "Shadows" database showing miniature military vehicle targets on the Army's former terrain-board at NVESD, Ft. Belvoir, VA. The scenes simulate desert and sparse forest locations, under simulated daylight or night-time flare lighting. Each scene contains 1 to 8 targets at a simulated distance of 1 to 3 km from the observer. Each image shows a 15° field of view, seen through a black-and-white camera that digitizes the image to 640 by 480 pixels with 256 possible gray levels. The eye-tracker data was collected with an *I-Scan* video-based system that collects eye position data at a rate of 60 Hz with an accuracy of about 0.5°. This raw data was post-processed by defining a fixation as at least six consecutive eye position samples falling within 0.5° of a running mean location (Hoffman, 1996).

The Fixation Probability Maps (FPM) generated from the set of 20 "Shadows" images were first visually compared to the eye-tracker data from the same set of images. For each of the 20 images, the FPM generated by our model was compared to a map generated directly from the eye-tracker data; each pair was found qualitatively very similar. Figure 7 below shows one example: above is the combined and smoothed eye-tracker data from 15 observers viewing the same input image (the image shown in Figure 2 above) and below is the FPM generated by our model from that same input image. Note that the most fixated objects on the lower left and near center are correctly predicted by the model.

The eye-tracker map in the figure below was generated as follows. First, we made the assumption that the 0.5° eye-tracker error is normally distributed with a reasonable variance. Instead of interpreting an eye-tracker fixation recording at (x,y) as a probability-of-fixation = 1.0 at (x,y) and 0.0 everywhere else, it should be interpreted as a 2-D normal distribution centered on (x,y). This distribution should be Gaussian with large variance (default value = 20) and mean chosen to make integrated volume = 1.0. Rather than adding the Gaussians one-at-the-time, a map was created to keep track of the number of recorded fixations per pixel. The location uncertainty due to eye-tracker error was then added by convolving the map with a Gaussian shape kernel as described above. (Note that using such a large kernel is made computationally reasonable by implementing the convolution as multiplication in the frequency domain.)
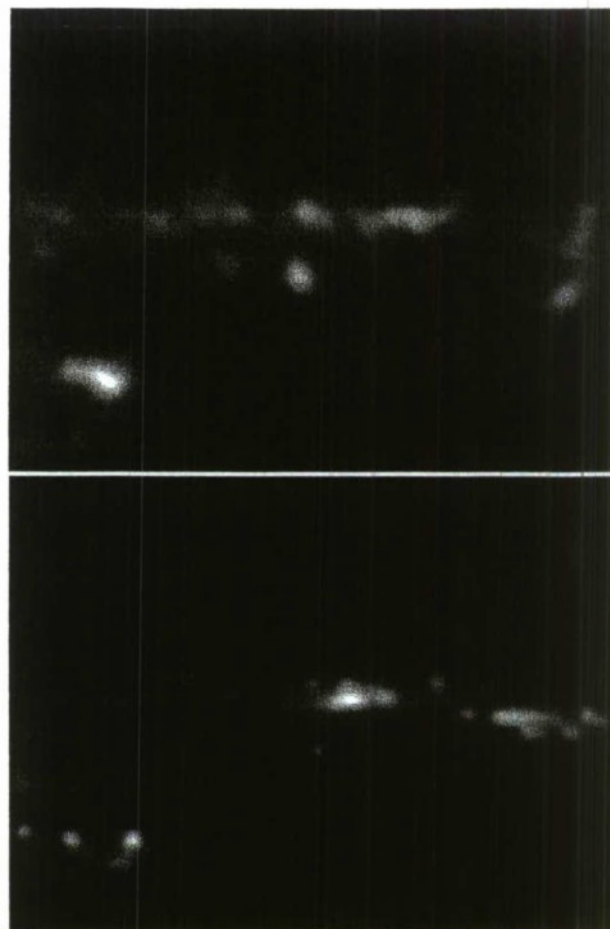
Figure 7: Eye-tracker data from 15 observers (above) and model-generated fixation probability map (below).

Quantitative comparisons were not as straight-forward because there is no such thing as "truth-data" for fixation locations and because of the large variance between individual observers. The standard approach to validating a human performance model would be to plot average human performance with error bars and then plot model performance on the same scale, analyzing how well the two graphs matched. This cannot be done for fixation data because there is no scalar measure of performance for fixation locations, there are no "wrong" or "right" fixation points. Our approach was to generate a "mean observer" map for each image, depicting the average fixations over all 15 observers (representing location uncertainty as described above) and use that as for all comparisons.

Figure 8 below shows the results per input image from two different comparisons with the mean observer: our model ($\blacktriangle$) and single observers ($\bullet$). Each comparison is made by computing the sum of pixel-by-pixel absolute differences between the mean observer map and the map in question, and then normalizing the result to the [0.0, 1.0] range. This Absolute Normalized measure of difference would be 0.0 for identical maps and 1.0 for

orthogonal maps. I.e., higher error values indicate more instances of a given location being fixated in one map and not in the other. The single-observer graph plots the average (and error bars) across 15 such comparisons, one per individual observer.

It is apparent from the figure that our model's FPM are quite similar to the mean observer. In fact, our model is *closer* (lower error) to the mean observer than most individual observers, for all but one image. This implies that treating our model like a 16th individual observer would not change the group statistics significantly.
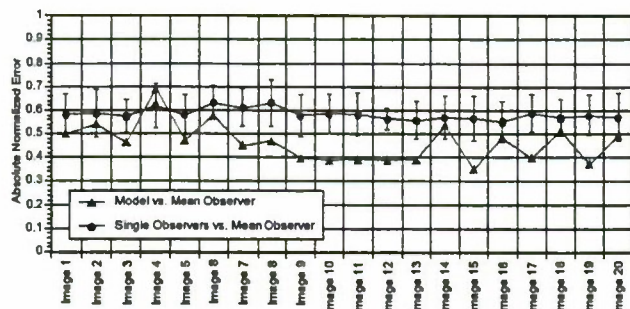
Figure 8: Absolute normalized error for model prediction (▲) and for single observers (●)

Another quantitative comparison that we performed was to compare two different versions of the model, with and without the Horizon Detector. The "Shadows" images had a visible horizon in every image and we therefore expected to see improved accuracy in FPM using the Horizon Detector. Figure 9 below confirms that expectation. The graph marked with (●) shows the results from the full model, including both the saliency and horizon bias inputs to each FPM (same as in Figure 8 above). The other graph (■ markers) shows the results from running the model with the Horizon Detector turned off, the error is increased for every image.
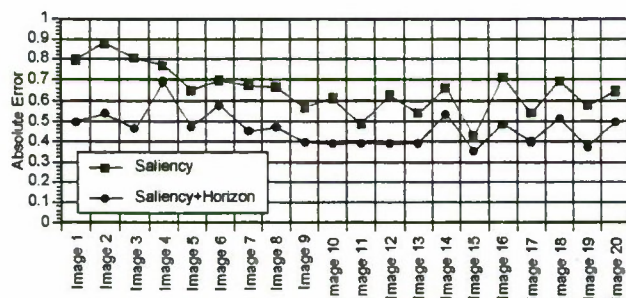
Figure 9: Comparing absolute normalized error for model predictions with and without horizon bias

We also qualitatively evaluated the output of the Fixation Selection component, by comparing it to maps of eye-tracker data. Figure 10 below shows the combined eye-tracker data for all 15 observers for the same input image

as before, with each fixation represented by a * overlaid upon the image. The circles (O) show the locations generated by the Fixation Selection component, using the full model with both saliency and horizon bias. Again, the most fixated objects on the lower left and near center are correctly predicted by the model.

**Figure 10: Eye-tracker data from 15 observers (*) and model-generated fixations (O)**

CONCLUSIONS

The major contribution of this project was the design and implementation of a computational model for predicting probability of fixation during search for hard-to-see targets. This model is image based: it uses *imagery* for input, rather than estimated parameter values characterizing critical factors such as clutter and target detectability. Consequently, it generates probability-of-fixation values that are primarily functions of image content, rather than functions of subjectively estimated parameters. The intended input domain is infrared or visible light imagery of distant vehicle targets in cluttered scenes.

Relatively more effort has been put into the search part of this model (predicting fixation locations) than the detection part (predicting target vs. non-target designation, given a fixation) reflecting our belief that the need for a mature search model is currently more important. Due to its modular design, the detection component of our model can easily be replaced by other detection models.

One of the characteristics of a modeling project of this scope is that some tasks turn out much harder than anticipated and others turn out to be unexpectedly simple. In the hope that our experience may be of use to future vision researchers, here are a few of the lessons we have learned:

- When modeling a function of the visual system, think in terms of images. Many researchers in the field of

target signatures are trained to work with zero or one dimensional highly abstract data, but the fundamental data type in the visual system is the 2-D retinotopic map. Abstracting the data to lower dimensions introduces abstract parameters that are very difficult to determine, even for expert users.

- Many aspects of eye-movements and visual function can be ignored in a visual search model. In particular, micro-saccades, fixation drift, and smooth eye-movements seem to be unimportant. If objects are far away then detection of object motion and stereo-range is not important (due to low velocity and disparity in retinal coordinates) and scan paths within objects are irrelevant (a single fixation covers the object).
- Fixation selection is a function of past fixations, but having the fixation probability map re-computed after each fixation is computationally prohibitive. A more reasonable solution is to use a static fixation probability map with a dynamic overlay, such as our fixation memory map.
- Partition your model according to trends in your experimental data, not according to "established approaches." We would never have thought of adding a Horizon Detector, which is a novel concept in visual search research, if we had not noticed a clear bias in fixation distributions towards the horizon.
- It is non-trivial to measure performance in the context of predicting fixation locations. Specific problems include lack of descriptive statistics that can tolerate slight spatial distortions (due to eye-tracker error) and also be useful for comparing model performance to the performance of one or more observers.
- Do not confuse designing an ATR system and a model of human vision. This may seem obvious, but it is an insidious problem because most engineers and researchers tend towards designing the system that performs best, instead of the system that performs no better than a human. Designing a model to make the *same mistakes* as a trained observer is the hardest part.

ACKNOWLEDGEMENTS

REFERENCES

Doll, T., McWhorter, S., Schmieder, D., Hetzler, M., Stewart, J., Wasilewski, A., Owens, W., Sheffer, A., Galloway, G., & Harbert, S. (1997). "Biologically-based vision simulation for target-background discrimination and camouflage/lo design", *Proceedings of SPIE*, Vol. 3062, pp. 231-242, Orlando, FL.

Gerhart, G., Meitzler, T., Sohn, E., Witus, G., Lindquist, G., & Freeling, J. R. (1995). "Early Vision Model for Target Detection", *Proceedings of SPIE*, Orlando, FL.

Hoffman, J. (1996). "Eye fixations during search of forest and desert terrain under simulated daylight and flare illumination conditions", NVESD Technical Report.

Itti, L., Koch, C., & Niebur, E. (1998). "A Model of Saliency-Based Visual Attention for Rapid Scene Analysis", *IEEE Trans. On Pattern Analysis and Machine Intelligence*, 20(11), 1254-1259.

Spillman, L., & Werner, J. (Eds.). (1990). *Visual Perception: The Neurophysiological Foundations*. San Diego: Academic Press.

Treisman, A., & Gelade, G. (1980). "A Feature-Integration Theory of Attention". *Cognitive Psychology*, 12, 97-136.

Wolfe, J. M. (1994). "Guided Search 2.0, A revised model of visual search". *Psychonomic Bulletin & Review*, 1(2), 202-238.

Wolfe, J. M. (1998). "Visual Search: A Review". In Pashler (Ed.), *Attention*, London, UK: University College London Press.
http://search.bwh.harvard.edu/RECENT%20PROJECTS/SEARCH%20REVIEW/reviewcontents.html

Analysis of TNO Search Data Using a 3-D Structure Contrast Metric

G. Witus
Turing Associates, Inc.
1392 Honey Run Drive
Ann Arbor, MI 48103 USA
E-mail: witusg@umich.edu

G. Gerhart
U.S. Army Tank-automotive and Armaments Command
AMSTA-TR-R / MS 263
Warren, MI 48397 USA
E-mail: gerhartg@tacom.army.mil

ABSTRACT

This paper presents an analysis of the TNO search data for the Search_2 imagery using a contrast metric based on the 3-D structure of the vehicle. The analysis employs the traditional P-infinity-times-negative-exponential model of search time distribution. P-infinity and mean search time are modeled as functions of the target signature. The signature metric is one over the product of vehicle size and contrast. The value of the metric is measured by the ability to account for variance in observed search performance.

Ground vehicles in natural lighting tend to have significant and systematic variation in luminance over the presented area. This arises, in large part, from the vehicle surfaces having different orientations and shadowing relative to the source of illumination and the position of the observer. These systematic differences create the appearance of a structured 3-D object. 3-D appearance is an important factor in search, figure-ground segregation and object recognition.

The 3-D structure contrast metric performs better than RSS contrast, and both perform dramatically better than the area-weighted average contrast. Target height performs better than either target area or square root of area. The signature metric accounts for over 80 percent of the variance in probability of detection and 75 percent of the variance in search time as measured in the TNO perception tests. When false alarm effects are discounted, the metric accounts for 89 percent of the variance in probability of detection and 95 percent of the variance in search time. The predictive power of the signature metric when it is calibrated to half the data and evaluated against the other half, is 90 percent of the explanatory power.

INTRODUCTION

Size and contrast have long been used to characterize the signature of simple targets in simple scenes for the purpose of analyzing search time and probability of detection. Size and contrast have been found to be good predictors of search and detection performance for stylized 2-dimensional targets, such as uniform disks and 4-bar patterns, against uniform backgrounds [Blackwell, 1943] [Ratches, et al., 1975].

Unfortunately, the standard area-weighted average contrast ratio has not proven to be a good predictor of search and target acquisition performance for complex targets in complex scenes. D'Agostino, et al., [1997] suggested a variety of possible modifications to the area-contrast metric to account for statistical luminance variation within the target and local surround. Peli [1996] concluded that the common measures of contrast are inadequate to explain detection performance for Gabor patches against uniform backgrounds, and suggested a computational contrast metric based on multi-scale band-pass filtering as an alternative.

Ground vehicles in natural lighting present non-uniform appearance when the surfaces of the vehicle are at different orientations with respect to the source of illumination and the observer (see fig. 1). The differences in shading between the adjacent surfaces reveal the 3-D structure. The appearance of common vehicles, from typical perspectives, under natural lighting is readily learned. This contributes to the perception of a 3-D object at a location, recognition of characteristic structural features, and classification as a potential vehicle.

This paper presents the initial results of exploratory research to develop a contrast metric based on the 3-D

vehicle structure, in natural lighting, relative to typical observer perspective. The objective of this research was to determine if a contrast metric could be defined based
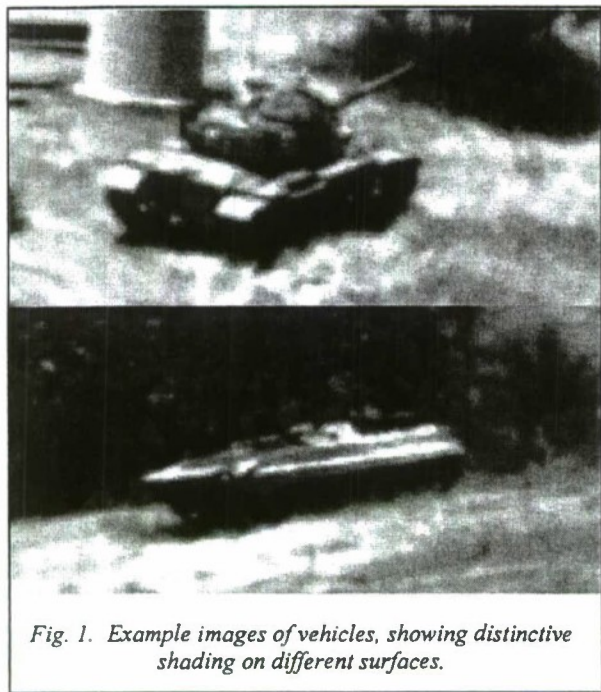

Fig. 1. Example images of vehicles, showing distinctive shading on different surfaces.

on the vehicle 3-D structure that would produce improved predictions of probability of detection and time to detect for military targets in natural backgrounds.

There is a substantial body of prior research suggesting that the perception of 3-D structure as a result of shape-from-shading is a significant factor in visual search and target acquisition. (Depth perception from visual parallax is insignificant at tactical ranges. For a stationary observer and a stationary target, shading and prior knowledge of vehicle appearance are the primary factors in 3-D shape perception.)

Marr [1982] coined the term "the 2½-D sketch" to refer to the perception of a 3-D structure from surface primitives. Sun and Perona [1996] showed that 3-D shading produced "pop-out" detection (i.e., response time independent of the number of distracters, indicative of pre-attentive parallel processing). They also showed that search became serial (time linear in the number of distracters) when the 3-D shading was removed even though the edge structure was retained. Tarr and Kersten [1998] concluded that the human visual system uses illumination angles to extract 3-D shape, and that illumination effects (including shadows) are modeled with respect to object shape, rather than simply encoded in terms of their effects in the image. Jonides and Gleitman [1972] and Mack and Rock [1998] both demonstrated that pre-attentive object recognition directs visual attention. Liu, Knill and Kersten [1995], and Liu and Kersten [1998] found that human efficiency exceeded 100 percent of an ideal 2-D observer for 3-D object classification. Moore and Cavanagh [1998]

demonstrated that perception of 3-D shape is possible from limited surface shading information, given familiarity with the 3-D object. Ullman [1996] has shown that observers use 3-D operations to recognize familiar objects presented in novel orientations. 3-D surface matching is also an approach being pursued for automatic object recognition systems designed to work in clutter with partially occluded targets (e.g., [Johnson and Hebert, 1998]).

MODELING APPROACH

Contrast, Size and Signature Metrics

This exploratory investigation employed a simplistic, low-resolution approach. If 3-D shading is a significant factor in search and target discrimination, then the effects should be apparent even though coarse analytic techniques were used. If coarse analysis does not reveal an effect, then the effect, if present, is probably not strong enough to be worth addressing in search and target acquisition models.

The conceptual 3-D vehicle model was based on the 3 cardinal surface orientations of a rectangular solid (vertical front, vertical side, and horizontal top). While military vehicles are not rectangular solids, the 3-region geometric model can be adapted with a little work. The projected view of a vehicle was divided into the following three regions (see figure 2):

1. Front/rear. The near-vertical, negatively sloped or self-shadowed portion of the front (or rear depending on the presented aspect). For armored vehicles this includes the lower glacis, front track/tire, and turret-chassis gap. For trucks, this includes the front grill, front of the cab, and front of the tires.

2. Side. The near vertical (e.g., within 10 degrees), negatively sloped or self-shadowed portion of the side, including the sides of the tracks or tires.

3. Top. All horizontal and near-horizontal surfaces up to a slope of 80 degrees. This includes all the small miscellaneous objects and protrusions on the vehicle. It includes the the upper glacis, top deck, roof, rear deck, turret armor. It also includes the sloped rear roof of the HMMWV.

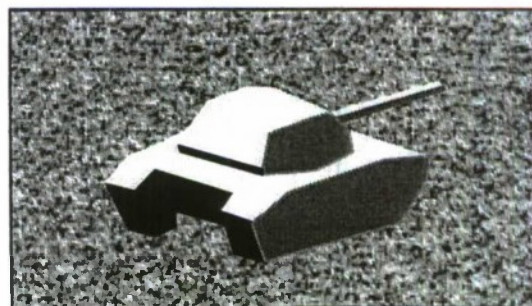These canonical surfaces are meant to identify the major

Fig. 2. Illustration of canonical front (rear), side, and top vehicle surfaces.

vehicle surfaces that typically have distinctive luminance resulting from different self-shadowing and angles relative to the observer and illumination. They address only the grossest level of 3-D structure. This level of resolution may be too coarse for modeling higher levels of target discrimination.

These regions also correspond to key structural features reported in field tests: darkly shadowed lower glacis, side profile, glint off roof or deck, lower grill, cab front, turret-chassis shadow. It is possible that the three surface orientations are significant because they correspond to important features for vehicle discrimination. It is also possible that the features are important because they reveal the 3-D structure.

The 44 Search_2 digital images [Toet, et al., 1998] were used in the demonstration analysis. All 44 images were used with no exceptions. The images were analyzed using Adobe Photoshop© to outline regions and compute gray-scale values. The local surround was taken to be a band surrounding the target with width equal to the target height.

The average gray-scale values for each region j, G_j, were converted to luminance values, L_j, via the calibration equation provided by Toet, et al., [1998]:

$$L_j = f(G_j) = 64.32 \left[(G_j - 18) / (G_j + 91.22) \right]^{2.3}$$

Since the calibration is a non-linear equation, a more accurate approach would have been to first convert pixel gray-scale to luminance, then compute the statistics.

The contrast for region j, C_j, is defined as the absolute value of the difference between the mean luminance of the region, L_j, and the mean luminance of the surround, L_{hkg}:

$$C_j = | L_j - L_{bkg} |$$

The vehicle contrast ratio metric, C_{veh}, is the area-weighted average of the contrasts of each of the three regions, divided by the luminance of the local background:

$$C_{veh} = \Sigma \, w_j \, C_j / L_{hkg}$$

where the weights, w_j, are the proportion of the presented vehicle area contributed by each region.

Two alternative contrast metrics were examined to provide a basis for relative comparison. These were the traditional area-weighted-average contrast [Ratches, et al., 1975] and the RSS contrast [D'Agostino, et al., 1997]. Both were computed by applying the non-linear gray-scale to luminance transform, f(), to statistics computed on the gray-scale images.

Signature metrics based on the area-weighted average contrast were uncorrelated with search performance (r^2 on the order of 0.3). Area weighted average contrast is not addressed in the remainder of this paper. This contrast metric was rejected.

The RSS contrast metric has been found to be an effective metric in other studies [D'Agostino, et al., 1997]. It is used as a reference for comparison with the 3-D structure contrast.

The RSS contrast ratio is the root-sum square of the target-background luminance difference and the target luminance standard deviation, divided by the mean background luminance:

$$RSS = \left[(\mu_{tgt} - \mu_{bkg})^2 + \sigma_{tgt}^2 \right]^{1/2} / \mu_{bkg}$$

For this comparison, the luminance standard deviation was estimated from the gray-scale mean and variance:

$$\sigma_{tgt} = \left[f([\mu_g^2 + \sigma_g^2]^{1/2})^2 - f(\mu_g)^2 \right]^{1/2}$$

where f() is the gray-scale-to-luminance calibration equation.

The signature metric, S_{veh}, used in the analysis is simply one divided by the product of the vehicle size measure, V_{veh}, and vehicle contrast measure, C_{veh}:

$$S_{veh} = 1 / (V_{veh} \, C_{veh})$$

The size metric in this analysis was the target minimum dimension, generally the vertical extent or height. Vehicle height was the measure of size used in the early Night Vision Laboratory target acquisition modeling [Ratches, et al., 1976]. Target height (vertical extent) was reported in the Search_2 documentation.

Two alternative size metrics have been proposed as alternatives to target minimum dimension: the vehicle presented area, and the square root of the presented area [D'Agostino, et al., 1997]. These size metrics were examined, but their performance was inferior to target height. Only analysis results using height are presented.

Search Model

The analysis employed the traditional search performance model that expresses probability distribution of detection over time as the product of a limiting probability of detection (P_{inf}) and a negative exponential distribution:

$$P_d(t) = P_{inf} * (1 - e^{-(t-\epsilon)/Td})$$

where ϵ is the minimum reaction time (nominally 1.0 seconds) and T_d is the mean time to detect given that a detection occurs [Washburn, 1981] [Ratches et al., 1975].

In the perception test subjects were given 60 seconds in which to search and respond [Toet, et al., 1998]. Toet reports the mean search time plus reaction time. To obtain the parameters of the search model, the effects of the 60-second response window and reaction time must be discounted. Assuming the negative exponential distribution of search time, given that a detection occurs, mean search time, discounting windowing and reaction time, can be computed from the reported mean search time, T_s:

$$T_d = (T_s - \varepsilon) / (1 - e^{-(60-\varepsilon)/T_s})$$

Toet et al. [1998] also reports the number of detections, N_d, false alarms, N_f, and misses, N_m. Probability of detection within 60 seconds can be calculated from this data:

$$P_d(60) = N_d / (N_d + N_f + N_m) = N_d / N$$

In the test image set, only one image had $P_d(60)$ less than 0.4, three images had $P_d(60)$ less than 0.5, five images had $P_d(60)$ less than 0.6, and 24 images had $P_d(60)$ greater than 0.95. Figure 3 shows the relative number of detection, false alarm and time-out (miss) responses in the perception test.

P_{inf} is computed from T_d and $P_d(60)$

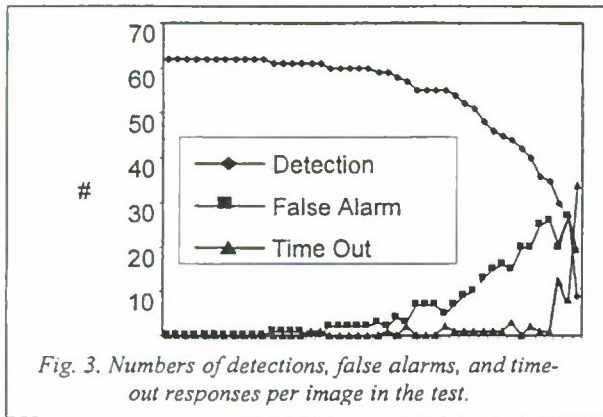$$P_{inf} = P_d(60) / (1 - e^{-(60-\varepsilon)/Td})$$

Fig. 3. Numbers of detections, false alarms, and time-out responses per image in the test.

The effects of the 60-second response window and the constant reaction time delay are minor. T_d is close to a linear function of T_s ($r^2=0.993$) and P_{inf} is close to a linear function of $P_d(60)$ ($r^2=0.992$) .

P-infinity Model

Given an unlimited search time, there are three possible outcomes: the observer can detect the target, false alarm, or conclude that there is no detectable target in the scene. Each of these is an absorbing state. As soon as the observer enters any one of these states, the search is over. Whenever a detection occurs, it is conditioned on having occurred before a false alarm and before the observer concludes that there is no target in the scene. In order for the conditional time to detect to have a negative exponential distribution, two criteria must be met:

(1) target detection, false alarms, and conclusion that no target is in the scene must be independent processes; and

(2) each of these processes must have a negative exponential distribution (albeit with different rates).

Under these conditions, the mean time to detect, conditioned on detection occurring first, is one over the combined response rate, R (R equals the sum of the individual rates of detection, R_d, false alarm, R_f, and concluding no detectable target is present, R_c):

$$T_d = 1 / R = 1 / (R_d + R_f + R_c)$$

P_{inf} is simply the ratio of the rate of true detection to the combined rates:

$$P_{inf} = R_d / (R_d + R_f + R_c)$$

These rates can be computed from the available data:

$$R_d = (N_d / N) R / (1 - e^{-59R})$$

$$R_f = (N_f / N) R / (1 - e^{-59R})$$

$$R_c = (N_m / N - e^{-59R}) R / (1 - e^{-59R})$$

The expected number of observers that decided no detectable target was present is equal to the probability of miss minus the probability that none of the three processes completed in the response window. Cinlar [1975] and Washburn [1981] provide details of the mathematics of competing Markov processes.

The perception test in which the Search_2 data was collected used 35 mm slides with targets present in every scene. The subjects knew that each scene contained a vehicle. The subjects also knew that they had only 60 seconds in which to search the scene. Under these conditions, the subjects would presumably continue searching for the full 60 seconds. Since they knew a target was present, they would not conclude no detectable target was present within the first 60 seconds. This implies that R_c should be zero.

This hypothesis is supported by the data. The mean value of R_c, computed over the 44 images, is 0.0005, and the standard deviation is 0.0018 (almost 4 times the mean). R_c is not statistically significantly different from zero, and even if it was, it is so small as to be insignificant for this analysis. The remainder of the paper disregards R_c.

Figures 4 and 5 show the distribution of the rate of target detection, Rt, and the rate of false alarm, Rf. Note that
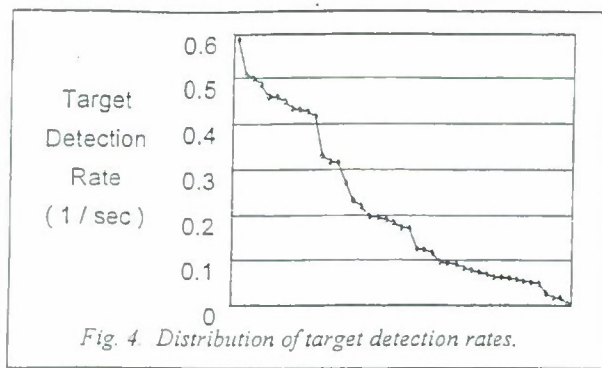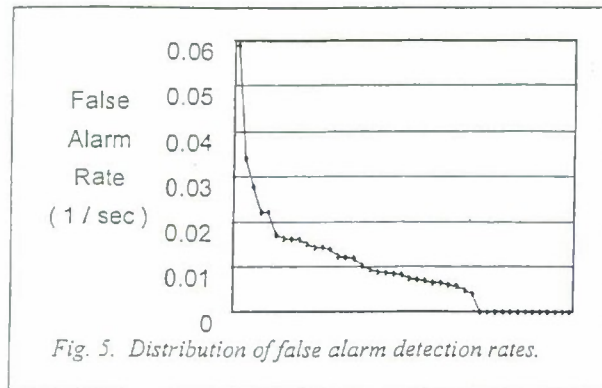
Fig. 4. Distribution of target detection rates.

Fig. 5. Distribution of false alarm detection rates.

the scales on the two graphs are an order of magnitude apart. Rt is greater than Rf for 43 of the 44 scenes.

The mean time to detect a target, given that a detection occurs, discounting the effect of false alarms is

$$T_t = 1/R_d$$

When the effects of false alarms are discounted, P_{inf} has value one. The probability that a detection occurs within 60 seconds, discounting the effect of false alarms can be computed directly from the response data on the numbers of detections and missed targets

$$P_d(60 \mid \text{no false alarms}) = N_d / (N_d + N_m)$$

When the effects of false alarms are discounted, P_{inf} has value one. The probability that a detection occurs within 60 seconds, discounting the effect of false alarms can also be estimated from the computed from the unconstrained mean time to detect a target without false alarms, Tt:

$$P_d(60 \mid \text{no false alarms}) = (1 - e^{-(60-\hat{\epsilon})R_d})$$

The root-mean-square (RMS) difference between these two estimates is 0.036, comparable to the sampling error in $P_d(60)$.

P_{inf} and T_d are modeled as simple linear functions of the signature metric. The model parameters (slope and intercept) are estimated from the data via linear regression. The related measures of search performance ($P_d(60)$, $P_d(60 \mid$ no false alarms), T_s and T_t) are modeled as functions of P_{inf} and T_d using the preceding search model equations.

ANALYSIS RESULTS

Gray-Scale Variance in the Vehicle Image

Partitioning the projection of the vehicle into the front, side and top regions accounted for 63 percent of the gray-scale variance over the entire target region. The area-weighted sum of the gray-scale variance within the three regions was 37 percent of the gray-scale variance over the entire target region.

This indicates that these vehicle regions account for a significant proportion of the gray-scale variance in images of ground vehicles. Sources of residual variance include small features defining local surfaces with different orientations and self-shadowing, paint patterns, shadows from trees falling on the vehicle, and patches of foreground obscuration.

The RSS contrast metric includes all variance over the vehicle, regardless of structural significance or spatial scale of the variation. The RSS contrast and 3-D structure contrast have a strong statistical linear relationship ($r^2 = 0.90$).

Sampling Error in the Search Performance Data

Sampling errors are inherent to any test procedure with a finite number of subjects. If the identical experiment were repeated with different subjects, the results would differ due to sampling error and the stochastic nature of search and detection.

$P_d(60)$ is estimated as the proportion of observers correctly detecting the vehicle. Assuming observer responses are independent, the sampling error has a binomial distribution. For a given image, the one-sigma sampling error in $P_d(60)$ is given by the following equation:

$$\sigma_{Pd} = [P_d(60) * (1 - P_d(60)) / N]^{1/2}$$

where N is the number of subjects (N = 62).

Figure 6 shows a plot of sampling error in $P_d(60)$ versus observed $P_d(60)$ for the 44 Search_2 images. The RMS sampling error in $P_d(60)$ over the entire Search_2 image set is 0.0363. The standard deviation in measured $P_d(60)$ over the entire image set is equal to 0.187. Sampling error explains 3.8 percent of the variance in $P_d(60)$ over the image set.

The long-run probability of detection, P_{inf}, was not measured directly, but was computed from measured data. This makes the effects of sampling error difficult to

Fig. 6. $P_d(60)$ vs. Sampling Error in $P_d(60)$.
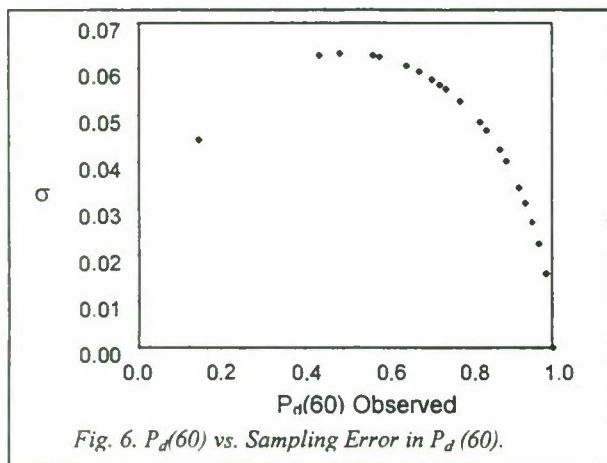
compute. However the effects of sampling error can be approximated by assuming the random variables were measured. The sampling error in P_{inf} is 0.036, explaining 4.4 percent of variance.

The probability of detection in 60 seconds absent the effects of false alarms, $P_d(60 \mid \text{no false alarms})$ is computed directly from the recorded data. The sampling error in $P_d(60 \mid \text{no false alarms})$ is 0.025, explaining 3.8 percent of variance.

The mean search time reported, T_s, is a constant reaction time ($\varepsilon = 1.0$ sec) plus a random variable with a negative exponential distribution truncated at 60-ε seconds. For this analysis the standard deviation is approximated by the standard deviation of a negative exponential random variable with the same mean (i.e., no truncation). The standard deviation of a negative exponential random variable is equal to the mean. Since T_s is computed using response time data only for subjects who detect th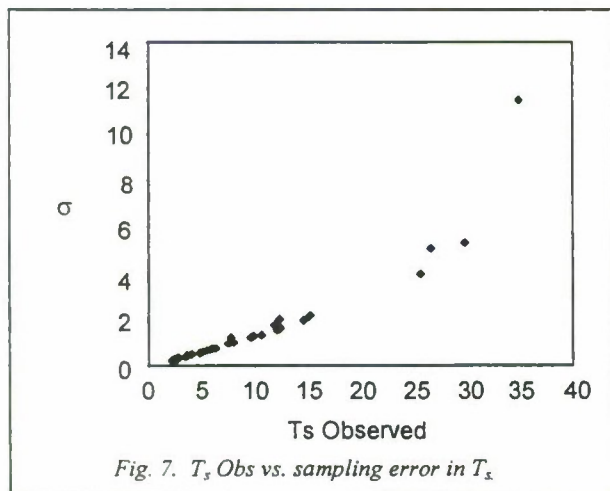e vehicle, for any given image the sampling error is equal to T_s divided by the square root of the number of subjects who correctly detected the vehicle:

$$\sigma_{Td} = (T_s - \alpha) / [N P_d(60)]^{1/2}$$

Figure 7 shows a plot of sampling error in T_s versus T_s

Fig. 7. T_s Obs vs. sampling error in T_s.

for the 44 Search_2 images. The RMS sampling error in T_s over the entire Search_2 image set is 2.4 seconds. The standard deviation in measured T_s over the entire image set is equal to 7.58 sec. Sampling error explains 10.3 percent of the variance in T_s over the image set.

The unconstrained mean time to detect, T_d, and the unconstrained mean time to detect in the absence of false alarms, T_t, were not measured directly, but were computed from measured variables. This makes the error due to sampling difficult to compute. The effects of sampling error can be approximated by assuming the random variables were measured. Both random variables have negative exponential distributions, so the 1-sigma sampling error is equal to the mean divided by the square root of the number of responding subjects. The estimated sampling error in T_d is 2.7 seconds, explaining 9.8 percent of variance. The estimated sampling error in T_t is 11.3 seconds, explaining 11.5 percent of variance.

Model Explanatory Power

The model has four free parameters that must be estimated from data: the slope and intercept of P_{inf} as a function of the signature metric, and the slope and intercept of T_d as a function of the signature metric.

The explanatory power of the model is measured by the percentage of variance in the observed search performance accounted for by the model. This is computed from the root-mean-square error between the model and observed data, and the variance in the observed data:

$$\%\text{Var} = 100 \, (1 - \text{RMS\_Error}^2 / \text{Observed\_Variance})$$

For a linear fit with parameters estimated via linear regression, the percentage of variance explained is equal to 100 times the Pearson correlation coefficient squared (r^2). Since the search model equations are non-linear, the percentage of variance accounted for is computed from the RMS error.

Figure 8 shows a scatterplot of the mean time to detect a target, given that the target is detected before a false alarm, but unconstrained by the 60 second time window of the experiment. The experimental value of T_d is computed from the measured search time. The model estimate of T_d is a linear function of the signature metric fit to the observed T_d.

Figure 9 shows a scatterplot of P_{inf} computed from the observed test data versus the linear function of the signature metric fit to the observed P_{inf} and truncated at one. Experimental values of Pinf are computed from T_d and raw response tallies.

271

Fig. 8. Unconstrained mean time to detect, t_d

Fig. 10. $P_d(60)$.

Fig. 9. $P_d(infinity) P_{inf}$.

Fig. 11. Mean search time, T_s.

Fig. 12. Mean Detection Time Absent False Alarms, T_t.
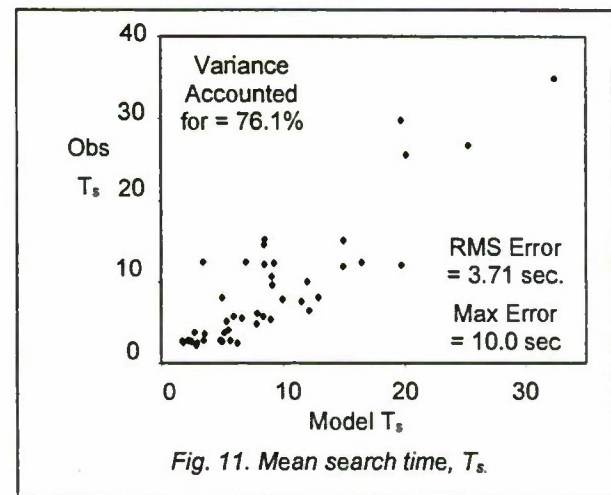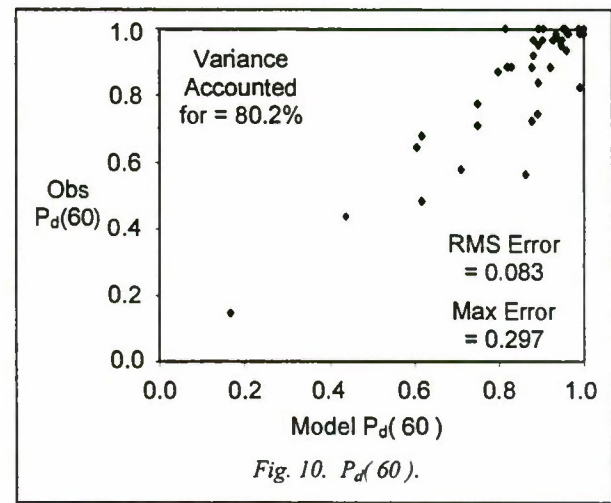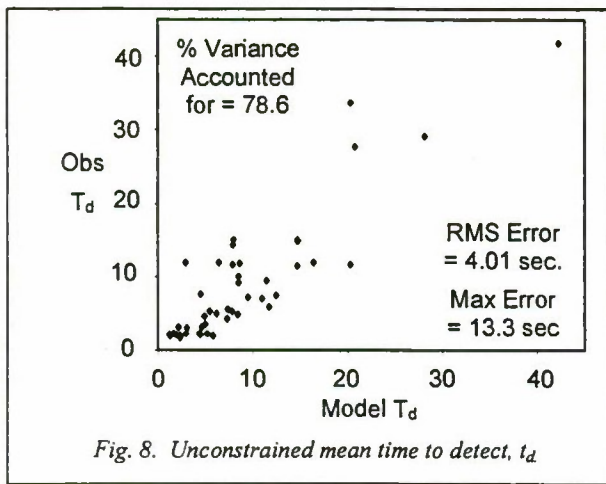
Figure 10 shows a scatterplot of the probability of target detection in 60 seconds computed directly from the tallies of observer detections, false alarms and misses, versus the model $P_d(60)$ computed from P_{inf} and T_d. The results are very similar to the P_{inf} results because, in most cases, the mean search time was much less than the 60 second response window.

Figure 11 shows a scatterplot of the mean search time measured in the experiment, versus the mean search time calculated by the model accounting for the effects of competing false alarms and the 60 second response window. These results resemble the results for unconstrained search time because, in most cases, the mean search time was much less than the 60 second response window.

Figure 12 shows a scatterplot of the mean time to detect a target without the requirement that the target detection occurs before a the first false alarm, T_t. It is the inverse of the rate of target detection. It is computed as T_d divided by P_{inf}. The experimental and model values of T_t are computed from the experimental and model values of values of T_d and P_{inf}. When the RSS contrast is used instead of the 3-D structure contrast, the percent of
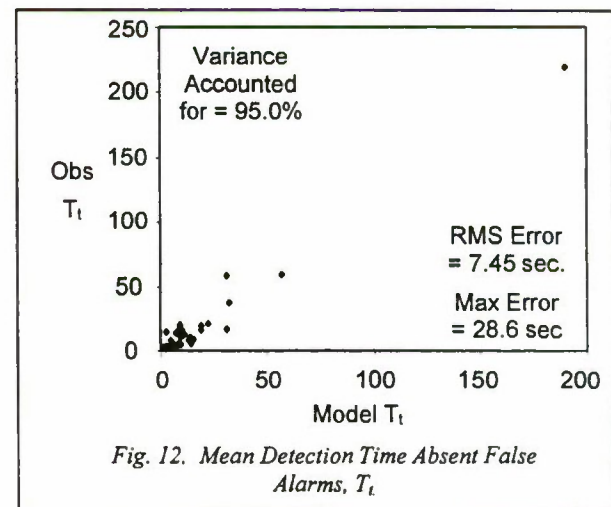
variance accounted for drops from 95 to 89 percent.

Many of the data points in figure 12 are clustered near the origin. The correspondence for low response time cases is seen more clearly when the logarithm of T_t is plotted (see figure 13). The logarithm operation is a non-linear

Fig. 13. Ln (mean detection time absent false alarms), Ln(T_t).

Table 1. Model Explanatory Power

Search Performance Measure	% Var	
	3-D	RSS
Unconstrained Time to Detect, T_d	78.6	77.8
P_{inf}	77.9	76.0
Search Time, Ts	76.1	75.2
$P_d(60)$	80.2	78.6
Detection Time Sans F.A., T_t	95.0	88.5
$P_d(60 \mid \text{No False Alarms})$	88.9	86.5
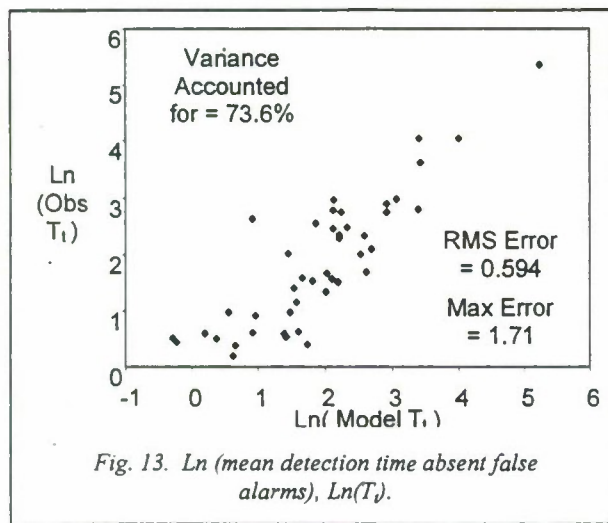
transformation, so the percent of variance accounted for is different.

Interestingly, the percent of variance accounted for by Ln(Model T_t) is equal to the percent of variance accounted for by linear regression of the signature metric directly on Ln(Observed T_t). When the RSS contrast is used instead of the 3-D structure contrast, the percent of variance accounted for drops from 74 to 66 percent.
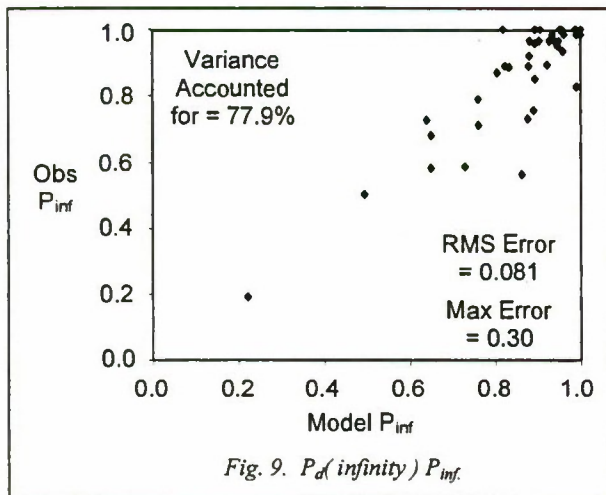
Figure 14 shows a scatterplot of the probability of target detection in 60 seconds, without competition from false targets, i.e., excluding false alarms. The experimental value is computed from the tallies of detections and misses. The model value is computed from T_t.

Table 2. Maximum Model Error

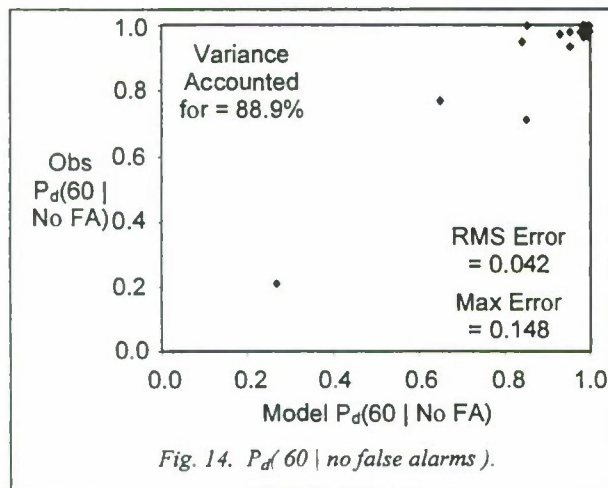| Search Performance Measure	Max Error	
	3-D	RSS
Unconstrained Time to Detect, T_d	13.3	11.9
P_{inf}	0.30	0.31
Search Time, Ts	10.0	9.0
$P_d(60)$	0.30	0.31
Detection Time Sans F.A., T_t	28.6	60.2
$P_d(60 \mid \text{No False Alarms})$	0.15	0.19

The results have several significant implications:

1. Signature metrics based on both the 3-D structure contrast metric and on the RSS contrast metric account for large proportions of the variance in search performance for this data set.

2. The 3-D structure contrast metric accounts for one to two percentage points more variance than the RSS contrast metric, except for the mean time to detect, absent false alarms where there is a 6.5 percentage points difference. This indicates that the 3-D structure contrast is a better measure observer response to the target signature. When the effects of false alarms are included, the additional variance due to this non-target source obscures the difference between the two contrast metrics.

Fig. 14. $P_d(60 \mid \text{no false alarms})$.

Tables 1 and 2 summarize the results of the comparison of the model to the data, and compare results obtained using the 3-D structure contrast metric with those obtained using the RSS contrast metric. Table 1 presents the percent of variance explained by the model. Table 2 shows the magnitude of the maximum model error.

3. The percentage of variance predicted by both metrics is significantly higher when the effects of false alarms are discounted. This is not surprising since the signature metrics do not measure potential false targets. The difference is greater for 3-D structure contrast than for the RSS contrast, further supporting

273

the claim that 3-D structure contrast is a better measure of the effects of the target signature.

4. The difference in the percent of variance predicted with and without the effects of false alarms indicates the magnitude of the contribution of false targets to search performance variance. By this measure, false targets account for over 15 percent of the variance in the mean time to detect a target, and almost 9 percent of the variance in the probability of target detection within 60 seconds.

5. The maximum error in P_d(60 | no false alarms) is significantly lower that the maximum probability error when the effects of false alarms are not excluded.

6. The magnitude of the maximum detection time sans false alarms is large. However this error occurs at the one hard-to-detect image, for which T_t was 218 seconds. The error, as a percentage of the time for that data point, is 13 percent for the 3-D structure contrast metric and 28 percent for the RSS metric.

Several excursions were conducted to assess alternative vehicle size metrics. When the signature metric was calculated using the square root of the presented area instead of the target height, the percent of variance predicted was approximately 12 percentage points lower for P_{inf} and 3 percentage points lower for T_d. When the presented area was used, the results were 15 percentage points lower for P_{inf}, and 6 percentage points lower for T_d.

Signature Metric Measurement Error

Measurement error occurs because the original images were blurred. The boundaries of the vehicles and regions in the vehicles were not sharply delineated. This affected both the measurement of target vertical extend and luminance. Not only was the location of the boundary uncertain, but pixels near the boundary contained a mix of target and background luminance, or a mix of the luminance between two regions.

Two separate estimates of the 3-D structure contrast ratio were made. Toet et al., [1998] provided one measurement of target height. A second independent measurement was made in this study. These measurements provided two pair of independent measures of the signature metric. Each independent pair of estimates produced one estimate of the measurement error in the signature metric.

The one-sigma measurement error in the signature metric over the Search_2 image set is 0.016. Since the model is linear, the measurement error in the predictions of P_{inf} and T_d are 0.016 times the magnitude of the slope (-2.313 and 92.11 respectively). This analysis yields a one-sigma measurement error in the predictions of 0.036 for P_{inf} and

1.473 for T_d respectively. The measurement errors in the predictions of P_{inf} and T_d are less than the sampling errors in the perception test estimates of P_{inf} and T_d (0.036 and 2.7 seconds respectively).

In combination, the variance due to sampling error and signature metric measurement error together are 9.1 percent of the P_{inf} variance predicted by the model, and 18.5 percent of the variance in T_d predicted by the model. The predictive power of the signature metric cannot be the result of spurious sampling and measurement errors.

The signature metric is one over the product of the vehicle 3-D structure contrast and the vehicle height. Two measurements of height and contrast were made, to obtain two pairs of independent measurements of the signature metric. The two correlations between the two pair of signature metric measurements were 0.986 and 0.979. The sample standard deviation for a pair of independent measurements is simply the difference between them. The error estimate for two pairs of independent measurements is the RMS of the two estimates of the sample standard deviation.

Figure 15 presents a plot of the signature metric measurement error sample standard deviation versus the signature metric value. The correlation is 0.91, suggest a strong linear relationship with slope equal to 0.15. As expected, the measurement error is larger for small, low-contrast vehicles than for large, high-contrast vehicles.
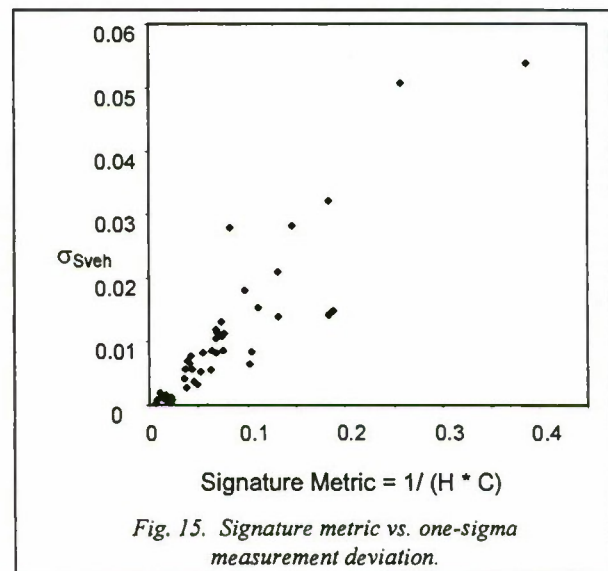
Fig. 15. Signature metric vs. one-sigma measurement deviation.

Accounting for Residual Variance

The model accounts for 75 to 80 percent of the variance in the experimental data when the effects of false alarms are included, and 90 to 95 percent of the variance when they are discounted. This suggests that for the TNO Search_2 data and test, false alarms account for 10 to 15 percent of the variance in probability of detection and search time, respectively.

Sampling error accounts for approximately 4 percent of the variance in probability of detection and 10 percent of the variance in search time. Together the target signature, false alarms, and sampling error are sufficient to account for all of the variance in the experimental data. (It is not possible simply to sum the percentage of variance explained by sampling error with the percentage of variance explained by the signature metric because of spurious correlation when the model parameters were estimated from the data).

The signature metric was calculated by applying the non-linear gray-scale-to-luminance transform to the mean and RMS gray-scale values. The correct method is to apply the gray-to-luminance transform to the image, then compute statistics. This approximation may account for some of the residual variance.

The contrast metric did not include any measure of color contrast or texture differences. The metric did not address chromatic, luminance or contrast adaptation, or spatial filtering. The metric did not address the effect of the position of the target in the scene, or its position relative to other features that might attract of inhibit attention to the target location. These factors may contribute to the model error, but the effect is likely to be small because the unexplained error is small.

There is no term that can be added to the signature metric to account for the residual variance. The prediction errors in P_{inf} and T_d are only weakly correlated ($r^2 = 0.14$). This suggests that once target signature effects are discounted, probability of detection and search time are sensitive to different processes and/or have non-linear relationships with image characteristics.

Individual Effects of Size and Contrast

One over the 3-D structure contrast metric was modestly correlated with perception test data ($r^2 = 0.7$ for T_d and 0.6 for P_{inf}). The percentages of variance explained for T_t and P_d(60 | no false alarms) were 51 and 54 percent. The RSS contrast metric had comparable correlation to T_d and P_{inf}, but accounted for 10 percentage points less of the variance in T_t and P_d(60 | no false alarms). The area weighted average contrast ratio had essentially no correlation with T_d or P_{inf}.

Target height, area and square root of area were only weakly correlated with T_d and P_{inf} (r^2 approximately equal to 0.4). Height had some correlation with T_t and P_d(60 | no false alarms) with r^2 on the order of 0.2. For area and square root of area, accounted for less than 10 percent of variance in T_t and and P_d(60 | no false alarms).

Height was less correlated with the 3-D structure contrast metric than it was with the RSS contrast metric ($r^2 = 0.3$ for the 3-D structure contrast metric, versus 0.4 for the RSS contrast metric).

These data indicate that height and 3-D structure contrast were largely independent dimensions, which individually were moderately correlated with search performance. Not surprisingly, their product was well-correlated with search performance. The same statements are true to a lesser extent for the RSS contrast metric.

Spurious Correlation and Predictive Power

When the same data are used to calibrate and evaluate the model, the percentage of variance accounted for is an accepted measure of explanatory (descriptive) power, but it is not truly a measure of the model's predictive power. In order to assess the model's predictive power, the model must be calibrated to one data set, then the prediction error evaluated for a separate, sequestered data set. This minimizes the effects of spurious correlation.
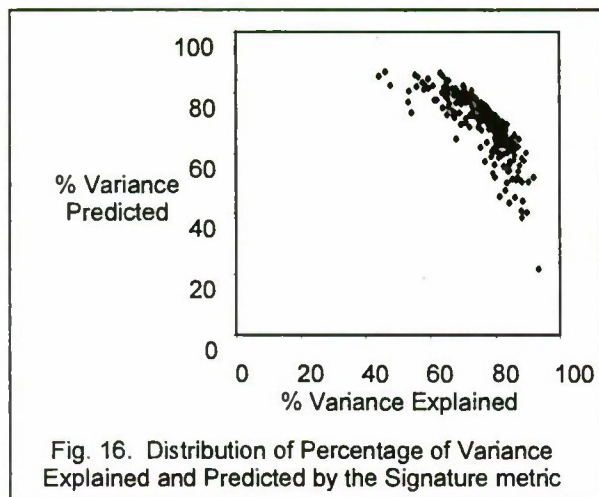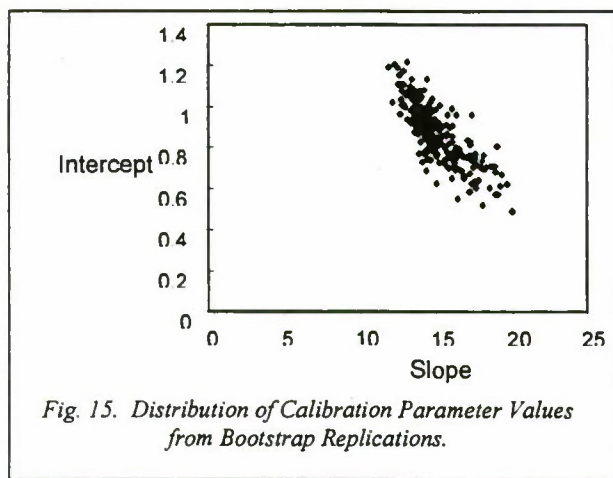
The Bootstrap statistical technique [Davison,. 1997] was used to evaluate the predictive power of the signature metric. The Bootstrap technique involves repeated random partitioning of the data into two disjoint sets: the calibration data set containing, and the validation data set. The model parameters are estimated from the calibration set, then the RMS prediction error is calculated from the sequestered validation data set. Each partition produces an estimate of the variance predicted by the model.

In this particular application of the Bootstrap technique, the calibration and validation data sets each contained half of the data points. Twenty-two calibration data points were used in the linear regression to estimate the model parameters, and 22 data points in the validation set were used to measure the RMS error and the percent of variance in the validation data set predicted by the model. Two-hundred-fifty-two (252) random partitions were generated to compute the Bootstrap statistics.

The Bootstrap analysis was applied to investigate the ability of the signature metric to predict the logarithm of the mean time to detect a target in the absence of false alarms, T_t. T_t was chosen as the dependent variable because it had the clearest causal relationship to the signature metric. The logarithm of Tt was used because of the uneven distributions of observed T_t and for the signature metric (see figures 12 and 13).

Figure 15 shows the distribution of slope and intercept from the 252 partitions. The slope had median value 14.6, equal to the slope when all 44 points are used in the regression (theBootstrap mean and variance are 15.0 and 1.6 respectively). The intercept has a median value of 0.89 compared to 0.90 when all 44 points are used in the regression (the Bootstrap mean and variance are 0.88 and 0.14 respectively).

Figure 16 shows the distribution of the percentage of variance in the calibration data sets versus the percentage of variance predicted in the validation data sets. The median percent of variance predicted in the validation

Fig. 15. Distribution of Calibration Parameter Values from Bootstrap Replications.

Fig. 16. Distribution of Percentage of Variance Explained and Predicted by the Signature metric

data sets is 72 percent (the mean and variance are 70 percent and 10 percentage points, respectively). The median percent of variance explained in the calibration data sets is 78 percent, compared to 76 percent when all 44 points are used in the regression (the Bootstrap mean and variance are 76 percent and 9 percentage points, respectively). On average (median and mean) the proportion of variance predicted in the validation data set is 92 percent of the proportion of variance accounted for in the calibration data set. This difference is due to spurious correlation, and indicates the difference between the explanatory and predictive power of the signature metric.

FINDINGS AND OBSERVATIONS

The traditional model of the distribution of search time was a useful model to analyze the experimental data.

With appropriate choice of definition of size and contrast, the simple signature metric equal to one over the product of size and contrast is a good fit to the observed data. It explains 75 to 80 percent of the variance in the test data, and 90 to 95 percent when the effects of false alarms are discounted.

The organization of the vehicles into three regions based on their orientation relative to the illumination and observer accounts for a significant portion of the gray-scale variance. Not surprisingly, the 3-D structure contrast metric and the RSS contrast metric are highly correlated and produce comparable results.

Nonetheless, the 3-D structure contrast metric is consistently superior to the RSS contrast metric, especially when the effects of false alarms are discounted. Variance due to false alarms obscures the difference in performance for the two contrast metrics. Both contrast metrics are far superior to the area weighted average contrast (which is no good at all).

Vehicle height is a better measure of target size, for use in product with a contrast metric, than either vehicle area or square root of vehicle area.

Responses to false targets, i.e., false alarms, account for a 10 to 15 percent of the search performance variance. Modeling the rate of false alarms as a function of the image properties has potential to improve search modeling.

There are a number of low-level and high-level visual phenomena not represented in this simple signature metric. Low-level factors include color contrast, chromatic and luminance adaptation, spatial filtering and contrast adaptation. Mid-level image processes include pre-whitening, edge detection and texture segregation. Beyond the vehicle structure, high-level (top-down) image properties include the location of the vehicles relative to terrain features that might attract attention or direct attention away, and position of the vehicle in the image.

These factors could account for the unexplained variance. However, they were not major contributors to search performance variance in the Search_2 image set. These factors could be more significant in other image sets containing greater variation on these dimensions.

The Search_2 vehicles do not present significant perceptible camouflage. Camouflage adds variance to the image. The RSS contrast metric will yield higher values for vehicles with camouflage than for vehicles without camouflage (assuming the same mean luminance), and thus will predict higher P_{inf} and lower T_d for camouflaged vehicles than for comparable non-camouflaged vehicles. The 3-D structure contrast metric is camouflage-neutral since it is based only on the mean luminance of different target regions and does not incorporate any higher-order statistics.

The Search_2 vehicles do, in some cases, have perceptible structures within the front, side and top regions. This increases P_{inf} and decreased T_d. These structures add variance to the image, which increases the value of the RSS contrast metric, which leads to higher

predicted P_{inf} and lower predicted T_d. The 3-D structure contrast metric is neutral with respect to structures within the three regions.

Neither the 3-D structure contrast metric nor the RSS contrast metric are able to distinguish modulation due to internal structure from modulation due to camouflage or foreground obscuration (e.g., brush or nets). More sophisticated signature analysis is needed to make this distinction.

ACKNOWLEDGEMENTS

REFERENCES

Blackwell, H. R., Contrast Thresholds of the Human Eye, *J. Opt. Soc.* 36: 624-43, 1943.

Cinlar, E., *Introduction to Stochastic Processes*, Prentice Hall, 1975.

Davison, A. C., and D. V. Hinkley, *Bootstrap Methods and Their Application*, New York: Cambridge University, 1997.

D'Augustino, J., W. Lawson and D. Wilson, Concepts for search and detection model improvements, *Proceedings of the SPIE* 3063: 14-22, 1997.

Johnson, A. E. and M. Hebert, Surface Matching for Object Recognition in Complex Three-Dimensional Scenes, *Image and Vision Computing* 16.9-10: 635-51, 1998.

Jonides, J. and H. Gleitman, A Conceptual Category Effect in Visual Search: O as a Letter or Digit, *Perception and Psychophysics* 12: 457-60, 1972.

Liu, Z. and D. Kersten, 2D Observers for Human 3D Object Recognition? *Vision Research* 38.15-16: 2507-19, 1998.

Liu, Z., D. C. Knill, and D. Kersten, Object Classification for Human and Ideal Observers, *Vision Research* 35.4: 549-68, 1995.

Mack, A. and I. Rock, *Inattentional Blindness*, Cambridge: MIT Press, 1998.

Marr, D., *Vision*, New York: W. H. Freeman & Co., 1982.

Moore, C. and P. Cavanagh, Recovery of 3D Volume from 2-Tone Images of Novel Objects, *Cognition* 67.1-2: 45-71, 1998.

Peli, E., In Search of a Contrast Metric: Matching the Perceived Contrast of Gabor Pathes at Different Phases and Bandwidths, *Vision Research* 37.23: 3217-24, 1997.

Ratches, J. A. et al., Night Vision Laboratory Static Performance Model for Thermal Viewing Systems, *R&D Technical Report ECOM-7043*, April, 1975.

Sun, J. Y. and P. Perona, Preattentive Perception of Elementary Three-Dimensional Shapes, *Vision Research* 36.16: 2515-29, 1996.

Tarr, M. J. and D. Kersten, Why the Visual Recognition System Might Encode The Effects of Illumination, *Vision Research* 38.15-16: 2259-75, 1998.

Toet, A. et al., A High-Resolution Image Data Set for Testing Search And Detection Models, *TNO-Report TM-98-A020*, Soesterberg, The Netherlands: TNO Human Factors Research Institute, 1998.

Ullman, S., *High-Level Vision: ObjectRecognition and Visual Cognition*, Cambridge: MIT, 1996.

Washburn, A. R., *Search and Detection*, Military Operations Research Society of America, 1981.

Classifier Performance Using Synthetic Data

Bruce A. Weber and Joseph A. Penn
Army Research Laboratory
AMSRL-SE-SE
2800 Powdermill Road
Adelphi, MD 20783

Abstract

In an effort to test and expand the usefulness of classifiers for U. S. Army applications, the U.S. Army Research Laboratory has begun an effort to create a large database of synthetic infrared images of vehicles. This database, once created, will be made available to qualified organizations for testing and evaluating their detector/classifier algorithms.

To date we have achieved an overall 50% recognition rate for a four vehicle database of real images, training our classifiers solely on synthetic data. It is expected that refinements to both the databases and the classifiers will improve this result.

In the discussion to follow, we will present the means to create the database, and the methods for and results of training and testing classifiers.

Introduction

Classifiers trained on real images of real vehicles have performed unreliably when tested on real images from other data collections. This is in part due to the large number of target classes and aspects, high-clutter background, sensor noise, and variations caused by translation, rotation, and scaling of vehicles. The problem is further complicated by inconsistencies in the signature of the vehicles, similarities between the signatures of different vehicles, limited training and testing data, obscured vehicles, and difficulty in using contextual information when it is available to the recognition system. These problems suggest that exploring data diversification and/or clutterless data may be useful in improving current classifier results.

To test these hypotheses we, at the U. S. Army Research Laboratory (ARL), are developing an image database of tens of thousands of synthetic infrared (IR) images of vehicles consistent with those in the COMANCHE IR database of real images. The COMANCHE database consists of approximately 15,000 vehicle images representing different image instances of 10 different vehicles, collected at three different geographical locations.

As a preliminary test, we created a small database of approximately 8000 synthetic vehicle images representing 4 of the 10 vehicles in different states of operation. For these same four vehicles there were approximately 6000 images in the COMANCHE database. Vehicle geometry and temperature distribution were derived using a BRLCAD/FRED/PRISM generator. Long wave infrared images (LWIR) were created using ARL's CREATION scene generator which alsoprovided the sensor model. The variations we considered were sensor to vehicle viewing aspect, time of day, time of year, longitude and latitude, direction of vehicle motion, and operating status.

This paper describes the synthetic image generator, the classifier, and the results of our training and testing this classifier using the synthetic and real image databases.

System Description

Figure 1 shows a schematic of the algorithmic methodology for the generation, training, and testing of the synthetic image database. Four programs were used to create the synthetic image database: BRL/CAD, FRED, PRISM, and CREATION.

The first three of these programs are used in series to create a temperature map of the vehicle. First, a CAD model of the vehicle is created in BRL/CAD by combining the intersection of multiple simple solid forms. These models contain descriptions of both the material and geometrical properties of these solids. At present, we have BRL/CAD models for 9 of the 10 vehicle types in the COMANCHE database.

After paring such a model down to a workable number of parts, FRED converts the solids into a series of multiply connected facets in the form of multisided polygonal plates. PRISM then combines the model's geometrical and material properties with weather, terrain, geographical location, time of day and year, and vehicle exercise information. This combination of data, when fed into the physical model of the heat production and transfer mechanisms, produces a temperature or radiance map for each facet at specific times throughout the duration of the exercise period.

Next, given these geometrical and time dependent temperature maps, CREATION renders tens of thousands of image instances of the vehicle as detected by an imaging sensor placed anywhere in the space surrounding the vehicle. Range plus sensor parameters, such as field of view (FOV) and optical bandpass can be varied to control image size and quality. Effects from atmospheric transmission, ground, sky, and solar radiance can also be added to adjust the image intensity. Finally accurate terrain maps from many real-world geographical locations can be chosen, which one then places the vehicle.
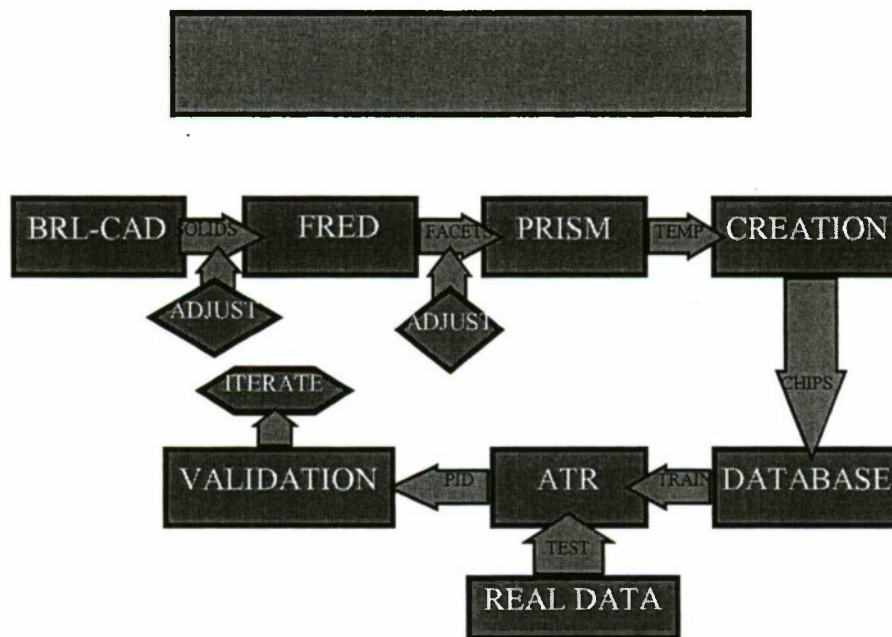
Figure 1. Methodology for the generation, training, and testing of the synthetic image database.

The classifier algorithm is used to train and test the synthetic images. Over the past three years two classifiers were developed at ARL for the training and testing of real IR images. One is a learning vector quantizer (LVQ) [1] and the other is a modular neural network (MNN) [2]. We chose to begin classification studies using the LVQ, since it is easy to interpret its performance, easy to analyze the results, and subsequently easier to determine how to improve the synthetic image database on the next iteration. Later we expect to use the MNN classifier. A description of the LVQ classifier follows in the classifier section.

Synthetic Data Preparation

For the purposes of this research, the validation of synthetic images of vehicles is defined to mean the extent to which such synthetic images can be used to train a classifier to recognize real images of vehicles. Thus, we determined early on that creating realistic synthetic images would be dependent on two real-world measurements: actual vehicle temperatures of specific vehicle parts and actual images of vehicles.

Based on other PRISM validation research, the BRL/CAD/FRED/PRISM trio of programs has been used successfully to create realistic temperatures accurate to within 2 to 4 ° C. But "create", we quickly found out, is a magical term, honed only by paying attention to detail and much experimentation. And the degree to which we succeed is defined by performance of our classifiers.

Given realistic temperatures for defined vehicle states, time of day, and weather conditions, "create" is an iterative process. But matching a particular temperature distribution for a particular set of initial conditions does not mean that one can predict accurately the future evolution of the temperature distribution.

In May of this year we gathered FRED/PRISM compatible files for three of the four vehicle types. The fourth was obtained in July. A preliminary synthetic set of images was generated and visually compared to corresponding real images. We

noted differences in signatures and adjusted the facet structure and heating mechanisms. This procedure was then repeated over several iterations, paying attention to real specific operational temperatures distributions where available.

Figure 2 shows such a comparison. On the left are real IR images of the four vehicles, and on the right are the synthetic examples. It is clear that although there are many similarities, there are also many differences. Some of the differences are due to time of day and vehicle orientation. Nevertheless, it is not clear whether such images will be effective for training the classifiers. Classifier training and testing must be performed now, before further adjustments are made.
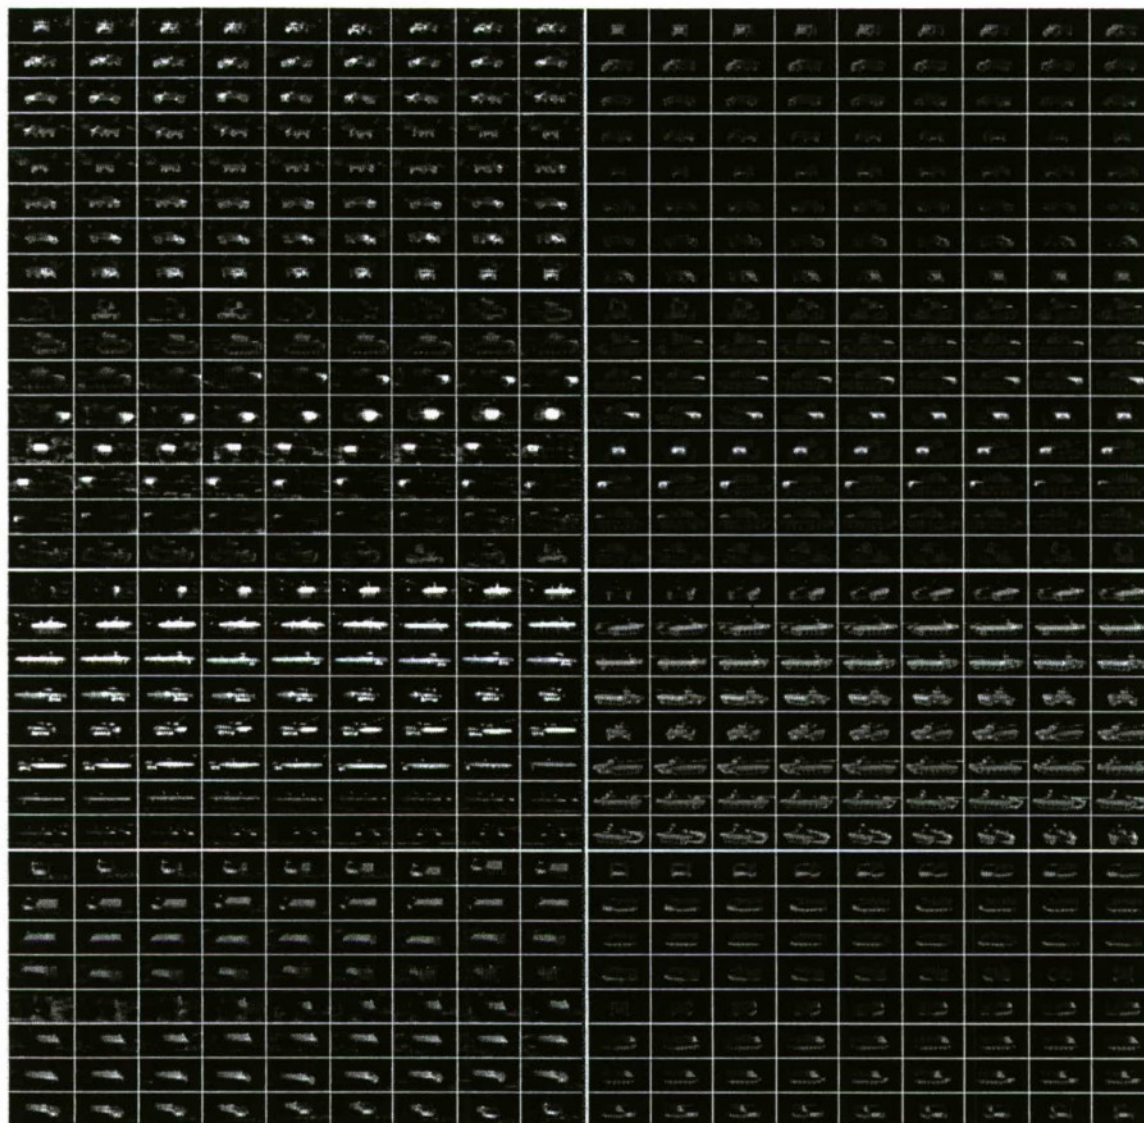
Figure 2. On the left are the real-world images of four vehicles. On the right are examples of the corresponding synthetic images. For each vehicle, there are 72 aspects representing every 5° step, beginning with zero degrees on the upper left, and with the vehicle rotated clockwise, as seen from above.

Classifier Description

The classifier used in this research, developed at ARL.[1], is described by its developers as a minimized mean squared error encoder. It consists of four stages as shown in the schematic in Figure 3

In the example presented in Figure 3, a set of eight aspect windows is used to extract a rectangular region of the image, based on the assumed aspect of the vehicle. Then the extracted region is enlarged into an area of a fixed size, decomposed by wavelet decomposition, and finally a dedicated vector quantizer (VQ) is applied for each sub-band within each aspect window.
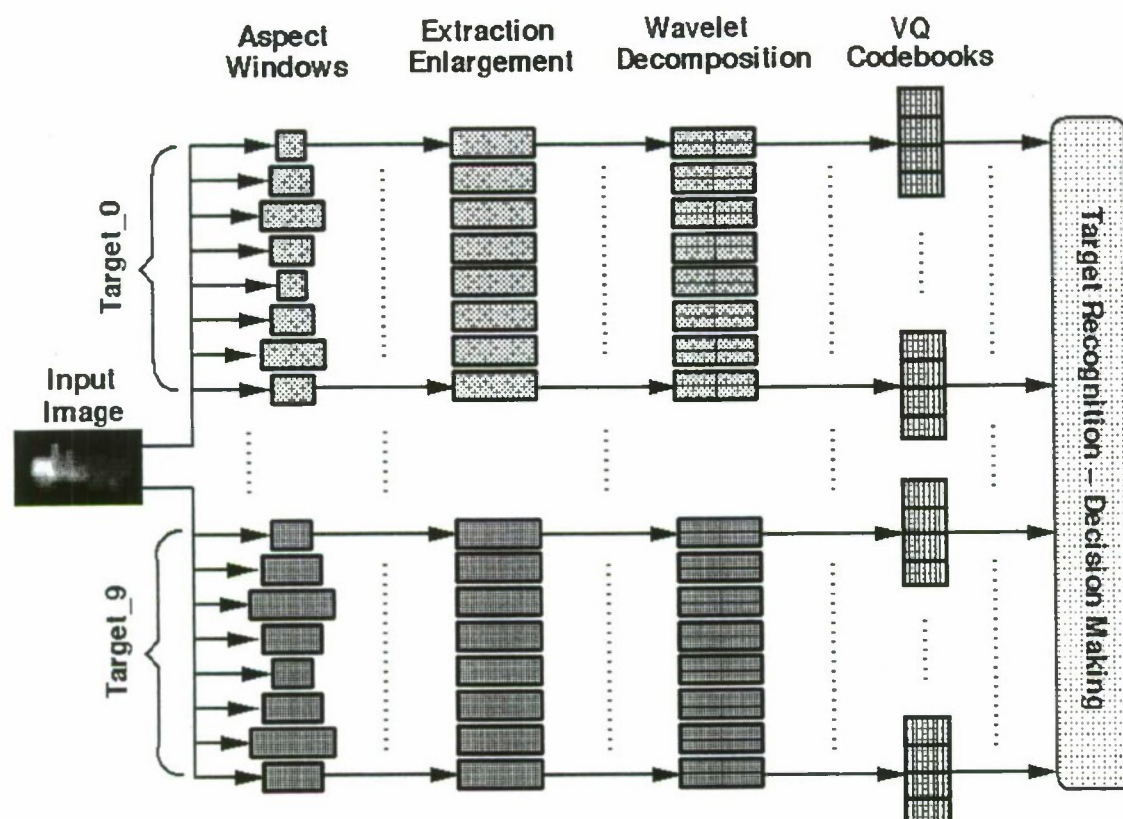
Figure 3. Schematic of LVQ algorithm [1].

From this a set of templates, or code vectors, is constructed for each sub-band of a particular target for a specific range of aspects. Each set of code vectors forms a codebook, representing the vehicle signature for a given sub-band of a particular vehicle for a specific range of aspects. A K-means algorithm trains the VQ independently by updating each code vector in each codebook with the average of all the data that are within a minimum distance, in terms of a Euclidean distance, to that code vector. The goal of this learning process is to minimize the average distortion, and thus capture the contextual similarities among the samples that belong to a particular sub-band of the intended vehicle and aspect window. The training stops when no more changes have been made to any of the code vectors.

For example, Figure 4 shows the 0° "0-0" sub-band codebook of the real image data on the left, and a sample synthetic data codebook for vehicle 3 is on the right. Note that having chosen 4 views (front, back, and two sides) for our codebooks, we see views spanning -45 to +45 degrees for this "0" degree codebook.

During the testing phase, each of the sub-bands of the extracted target area is represented by a similarity measure that compares the given sub-band with the best-matching code vector from the corresponding codebook. A commonly used similarity measure, called the mean squared error (MSE), is used to compare a code vector and the sub-band being tested. The input image class of the group of sub-band codebooks produces the smallest MSE.

283
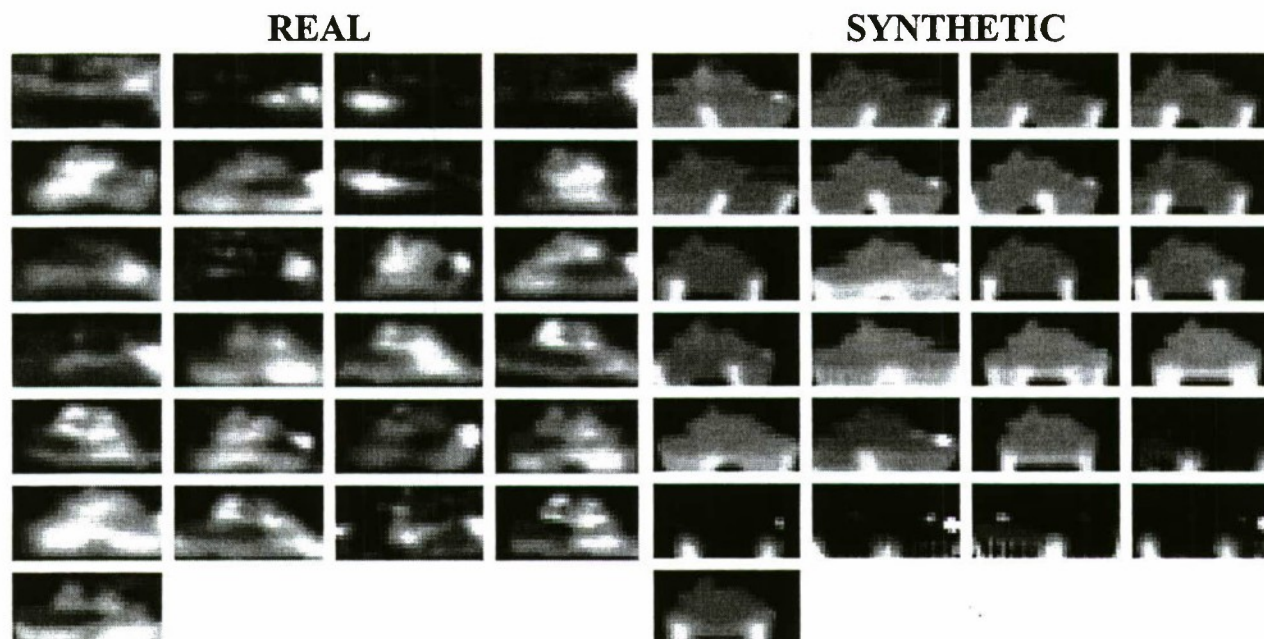
Figure 4. The 0° "0-0" sub-band codebooks for vehicle 3.

For simplicity, we are using four aspect windows and four wavelet sub-bands for each of the four vehicles. In all, a total of 64 possible paths must be checked for each classification.

Results

Table 2 shows the results of training on the 5156 real image database and testing on the 8000 synthetic image database. . In this case no real images were used to train the classifier. For this table, and for similar tables to follow, the diagonal elements beginning in the upper left and descending toward the lower right indicate the percentage of the vehicles tested that were correctly identified.

Table 3 shows the results of training on the 8000 synthetic image database and testing on the 5156 real image database.

At present, the overall recognition rate for training the VQ classifier on synthetic data alone and testing on real data is 50%, whereas for training on the real data and testing on the synthetic data the result is 68%.

Based on our preliminary results, we believe that continued refinement of our synthetic database will further improve our probability of correct identification rates.

Table 2. A confusion matrix showing the probability for correct synthetic image vehicle identification having trained on only real images.

Probability of Correct Vehicle Identification (%)				
Predicted/Actual	Vehicle 1	Vehicle 2	Vehicle 3	Vehicle 4
Vehicle 1	79.76	5.11	17.21	0.25
Vehicle 2	0.15	91.32	51.93	0.00
Vehicle 3	11.81	2.33	30.61	27.83

Vehicle 4	8.28	1.24	0.25	71.92

Table 3. A confusion matrix showing the probability of correct real image vehicle identification having trained on only synthetic images.

Probability of Correct Vehicle Identification (%)				
Predicted/Actual	Vehicle 1	Vehicle 2	Vehicle 3	Vehicle 4
Vehicle 1	51.16	13.18	12.13	17.05
Vehicle 2	39.52	72.60	49.39	20.37
Vehicle 3	3.96	6.32	26.22	4.81
Vehicle 4	5.36	7.91	12.26	57.78

Conclusions

We began this project hopeful of attaining quick results to support further investigation into the application of synthetic images for the training and testing of vehicle classifiers. Each iteration of our database provides more realistic and more variable synthetic images. The performance of our database has steadily improved, and we expect continued improvement.

We have learned to pay careful attention to the details of those regions of the vehicle in which thermal changes are frequent and substantial. Most models needed to have a few large regions cut into smaller regions so that realistic thermal diffusion could be properly represented.

Validated thermal data for each vehicle are needed to help guide our modeling. These data may exist, but at present they are not available to us. Additional data representing the thermal evolution of signatures over time are also not available to us or are nonexistent. Any help from the research community in obtaining such data would be greatly appreciated.

The ultimate purpose of our research to derive data capable to aiding the training and testing of classifiers. It is the intent of ARL, when this work is completed, to make the entire 10-vehicle synthetic database available to qualified organizations.

Acknowledgments

The authors are grateful to Lipchen A. Chan for his discussions and great patience while we were coming up to speed in the use of the LVQ algorithm. We would also like to thank Dan Rinald of the National Ground Intelligence Center for taking two days out of his working life to give us a tutorial into the proper application of BRL/CAD, FRED, and PRISM. Finally, we would like to thank Eddie Jacobs of Night Vision Electro-Optic Laboratory for supplying a PRISM-ready model for one of the four vehicles in our database.

References

[1] Lipchen A. Chan, Nasser M. Nasrabadi, and Vincent Mirelli, 1996, Wavelet-Based Learning Vector Quantization For Automatic Target Recognition, SPIE Proc. Aerospace/Defense Sensing and Controls, Vol. 2755, pp. 82-93, Orlando, April 1996.
[2] Lin-Cheng Wang, Sandor Z. Der, and Nasser M. Nasrabadi, 1998, Automatic Target Recognition Using A Feature-Decomposition and Data-Decomposition Modular Neural Network, IEEE Transactions on Image Processing, Vol. 7, No. 8, August 1998.

Target Classification and Support Vector Machines

Robert E. Karlsen, David J. Gorsich and Grant R. Gerhart
U.S. Army Tank-automotive and Armaments Command
Warren, MI 48397-5000

ABSTRACT

The area of automatic target classification has been a difficult problem for many years. Many approaches involve extracting information from the imagery through a variety of statistical filtering and sampling techniques, resulting in a reduced dimension feature vector that is the input for a learning algorithm. In this paper, we introduce the Support Vector Machine (SVM) algorithm, which is a wide margin classifier that can provide reasonable results for sparse data sets and whose training speed can be nearly independent of feature vector size. Therefore, we can avoid the feature extraction step and process the images directly. The SVM algorithm has the additional features that there are few parameters to adjust and the solutions are unique for a given training set. We applied SVM to vehicle classification and character recognition problems and compared the results to standard neural network approaches. It was found that the SVM algorithm gave equivalent or higher correct classification results compared to the neural networks.

INTRODUCTION

The problem that we address in this paper is image classification. Specifically, we consider algorithms that can take as input a digital image and classify it according to some criterion. Often this is a three-step process, consisting of pre-processing, feature extraction, and decision algorithm. The pre-processing is used to remove redundant information or to transform the image to a space where the objects are more easily classified. While we have used the multiresolution approach, implemented through the fast wavelet transform, in previous applications,[1,2] for this work we have performed little pre-processing.

The feature extraction step is often employed to reduce the dimensionality of the problem. Examples of features include peaks in the Fourier spectrum, statistical measures of edge densities, multiresolution energies or a histogram of gray levels. In many cases, the type of problem dictates a certain choice for the feature vector. However, in image analysis, the correct features are not always well defined or easily extracted and much time and effort can be expended searching for an appropriate feature space. In the current application, we avoid feature extraction and work directly with the image.

The final step in many systems is a decision module, which takes as input the lower dimension feature vector and outputs the classification. Often the decision algorithm contains a learning module where a sufficiently large number of sample images, with their associated classification, are presented to the algorithm. A variety of internal parameters are automatically adjusted in order to satisfy a minimum error criterion so as to achieve optimal classification results. The typical learning algorithm is a neural network (fuzzy logic systems generally have an embedded neural network for automated learning). The reason that feature vector selection is so important for neural networks (specifically Multi-Layer Perceptron neural networks) is that the complexity of the network scales with the size of the feature vector. In fact, the number of free parameters that must be determined is proportional to the size of the feature vector and is often many times larger. One would expect that a large number of training samples is required in order to constrain the error minimization sufficiently and provide good generalization.

The Support Vector Machine (SVM) algorithm[3-5] avoids many of these difficulties. Here, only the scalar product between feature vectors enters the problem. Therefore, the length of the feature vector has little effect on the computational complexity of the algorithm. By design, SVM is a wide margin classifier and can give reasonable results even for sparse training sets, where the number of samples may be less than the size of the feature vector. The algorithm also provides unique results for a given training set and can be made resistant to outliers by adjusting a cost parameter. Previous results have shown
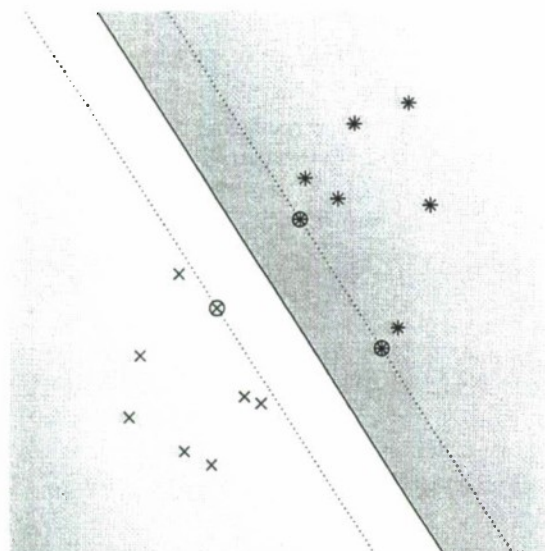
Fig. 1: Linearly separable data.

that the SVM algorithm is comparable to the best neural networks.[6] Since the algorithm is fairly new, the hope is that with further research it will, on average, surpass neural networks in ease of use and classification performance. A big advantage of the SVM algorithm is that it has few adjustable parameters that need to be tweaked in order to achieve good results.

The paper begins by providing a brief tutorial on Support Vector Machines and outlining derivations of some of the important elements. A fast algorithm for SVM training is also introduced. The SVM algorithm is then applied to a collection of images of military vehicles and the resulting classification results are compared with a Multi-Layer Perceptron (MLP) neural network and with a Radial Basis Function (RBF) neural network. The fast algorithm is also compared with the MLP on a larger character recognition task. We conclude by summarizing our results and indicating the future direction of our research.

SUPPORT VECTOR MACHINES

Linearly separable data

The SVM algorithm[3-5] is based on finding the pair of parallel hyperplanes that separates the data into two classes and has the largest perpendicular distance between them. We conjecture that this will provide a good approximation to the "best" separating hyperplane, where by "best" we mean the hyperplane that will, on average, give the lowest classification error when new data is used.

To introduce the Support Vector Machine (SVM), we first consider a linearly separable problem; i.e. the data can be separated completely by a hyperplane. Figure 1 shows an example in two dimensions, where the hyperplane is a line. Each data point is described by a feature vector, $\bar{x}$, and a truth value, y, that can take the values of $+1$ or -1, depending on the class. The two hyperplanes are required to pass through at least one point of each class and there can be no points between them. The boundary between the classes is then defined to be a third parallel hyperplane that is halfway between these two. The data points that the outer hyperplanes pass through, which are circled in Fig. 1, are called the support vectors, the meaning of which will be explained later. One of the hyperplanes consists of those points x which satisfy,

$$\bar{w} \cdot \bar{x} + b = +1, \tag{1a}$$

while the other hyperplane contains those points which obey,

$$\bar{w} \cdot \bar{x} + b = -1, \tag{1b}$$

with Eq. (1a) going through at least one point of class $y=+1$ and Eq. (1b) going through at least one point of class $y=-1$. The constants $\bar{w}$ and b define the hyperplanes, where $\bar{w}$ is a vector that is normal to the hyperplanes and $-b/\|\bar{w}\|$ is the perpendicular distance from the origin to the middle hyperplane. The RHS of Eq. (1a) will be greater than or equal to $+1$ for all points of class $y=+1$ and the RHS of Eq. (1b) will be less than or equal to -1 for all points of class $y=-1$. These can be combined into the following constraint on all the data points,

$$y_i\left(\bar{w} \cdot \bar{x}_i + b\right) - 1 \geq 0 \quad (1 \leq i \leq n). \tag{2}$$

The perpendicular distance between the two outer hyperplanes is equal to $2/\|\bar{w}\|$. Therefore, finding the hyperplanes with the largest margin reduces to computing values for $\bar{w}$ and b that minimize $\|\bar{w}\|^2$, subject to the constraint in Eq. (2).

A standard procedure for handling optimization problems with constraints is given by the Lagrangian formalism.[7-9] The constraints are taken into account by adding multiples of the constraint equations to the objective function, which in this case, results in the following primal Lagrangian,[3-5]

$$L_P = \frac{1}{2}\|\bar{w}\|^2 - \sum_i \alpha_i y_i\left(\bar{w} \cdot \bar{x}_i + b\right) + \sum_i \alpha_i, \tag{3}$$

287

where α_i are the Lagrange multipliers associated with each of the constraints in Eq. (2). Because the constraints are inequalities, bounded from below, the Lagrange multipliers are required to be non-negative.[7-9]

Setting the derivative of the Lagrangian in Eq. (3) with respect to $\bar{w}$ and b (the primal variables) equal to zero, leads to the following expressions,

$$\bar{w} = \sum_i \alpha_i y_i \bar{x}_i, \tag{4a}$$

$$\sum_i \alpha_i y_i = 0. \tag{4b}$$

Inserting Eqs. (4a) and (4b) into (3), results in the dual Lagrangian,

$$L_D = \sum_i \alpha_i - \frac{1}{2} \sum_{i,j} \alpha_i \alpha_j y_i y_j \bar{x}_i \cdot \bar{x}_j. \tag{5}$$

The problem is now reduced to finding the Lagrange multipliers (the dual variables) that maximize Eq. (5) and satisfy both the non-negativity constraints and the constraints of Eq. (4b). From the theory of constrained optimization ones finds that the only constraints that matter are those that actually constrain the minimization of the objective function. The other (inactive) constraints can be discarded without changing the optimal point. It is consistent therefore to set the Lagrange multipliers for the inactive constraints to zero.[8,9] This condition can be summarized as:

$$\alpha_i \left(y_i \left(\bar{w} \cdot \bar{x}_i + b \right) - 1 \right) = 0, \tag{6}$$
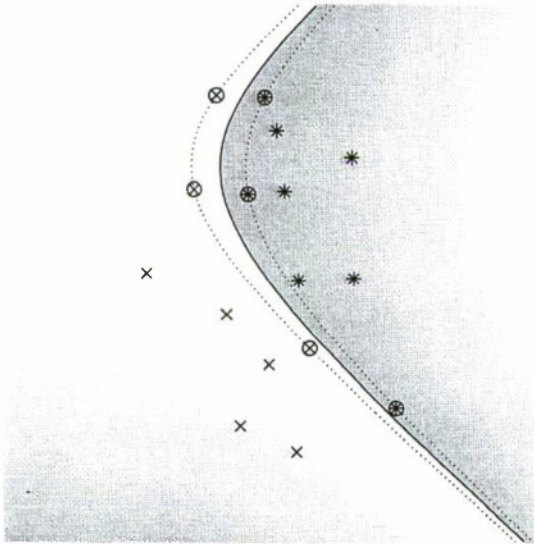

Fig. 2: Nonlinearly separable data.

which means that those data points with non-zero Lagrange multipliers (and hence are active constraints) will lie on the outer hyperplanes. These data points are called the support vectors and they are the points that determine the position of the hyperplanes. One can move the other points around the feature space or remove them entirely and the solution will not change, provided one does not move a point across one of the outer hyperplanes.

One can solve Eq. (5) using any quadratic programming solver, although different solvers perform better on different types of problems.[5,7-9] Solving the quadratic programming (QP) problem efficiently is actually one of the most difficult parts of SVM and there exists many numerical QP solvers that are readily available.

Once the Lagrange multipliers are known, the solution for $\bar{w}$ is given by Eq. (4a), where the sum can be restricted to the support vectors, since they are the only ones with non-zero α. One can find b from Eq. (6) using any of the support vectors, although one generally averages over all the support vectors for better accuracy. Once these constants are known, the classification of an unknown vector, $\bar{v}$, is given by the sign of,

$$b + \sum_i \alpha_i y_i \bar{x}_i \cdot \bar{v} \quad (1 \le i \le n_{sv}), \tag{7}$$

where the sum is over the support vectors.

Nonlinearly separable data

Now suppose that the "best" boundary between the data is nonlinear. An example of this situation is shown in Fig. 2, using a non-homogeneous quadratic kernel. One cannot separate the two classes with a straight line. The structure of the SVM equations allows a simple solution to this situation. Map the data, through a nonlinear transformation ϕ to a different space where the data can be separated with a hyperplane. This results in the Lagrangian in Eq. (5) being transformed to,[3-5]

$$L_D = \sum_i \alpha_i - \frac{1}{2} \sum_{i,j} \alpha_i \alpha_j y_i y_j \phi(\bar{x}_i) \cdot \phi(\bar{x}_j). \tag{8}$$

Since Eq. (8) depends only on scalar products between the transformed feature vectors, one can replace the scalar product with a kernel function,

$$K(\bar{x}, \bar{y}) = \phi(\bar{x}) \cdot \phi(\bar{y}), \tag{9}$$

and never need to compute the transformation ϕ explicitly. Equation (7) then becomes,

$$b + \sum_i \alpha_i y_i K(\bar{x}_i, \bar{v}) \quad (1 \le i \le n_{sv}), \tag{10}$$

with the test feature vector now inside the summation over the support vectors.

Since one is still solving the linear problem, just in a different space, the computational overhead is essentially the same. The solution and parameters for the hyperplane are in the higher dimensional space and when one transforms back to the original space the boundary becomes nonlinear. However, there is, in general, no way to analytically invert the solutions for $\bar{w}$ and b. Hence, one must use Eq. (10) to classify feature vectors. The advantage to using the kernel approach is that the higher dimensional (or embedding) space is essentially hidden from the user. One, in fact, never needs to know the function ϕ. It could even be of infinite dimension. Determining the best kernel for a given problem is a subject of active research. We have used a non-homogenous quadratic function in our numerical work:

$$K(\bar{x}, \bar{y}) = (\bar{x} \cdot \bar{y} + 1)^2. \tag{11}$$

Non-separable data

A potential problem can occur when the data is not separable using a given kernel. This is demonstrated in Fig. 3, where, due to an outlier, the data cannot be separated with a linear kernel. Then the assumptions leading 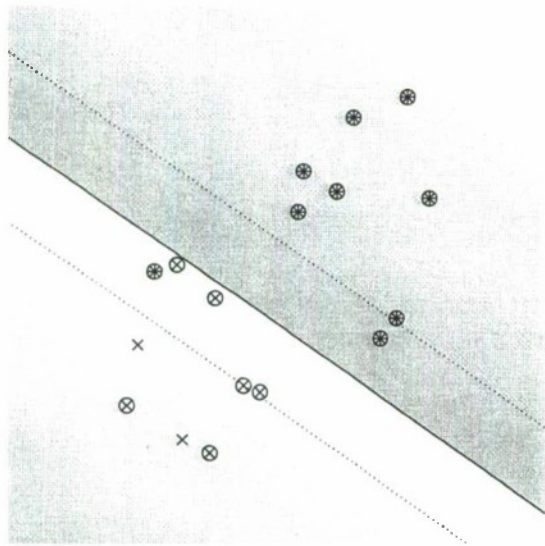to Eq. (1) no longer hold. Although the preceding SVM algorithm can provide a reasonable solution in these cases, many times the separating hyperplane is not one that would be considered the "best" solution. This is due to the outliers being given more weight than the other data points. To avoid this, positive slack variables δ are introduced that measure the deviation of the outliers from the optimal separating hyperplanes. The constraint equation (2) then becomes,[3-5]

$$y_i(\bar{w} \cdot \bar{x}_i + b) - 1 + \delta_i \ge 0. \tag{12}$$

A convenient way to minimize the total amount of the outlier error is to add an appropriate term to the Lagrangian multiplied by a constant C. Choosing the error term to be the sum of the deviations leads to the optimization solution being independent of the slack variables and their associated Lagrange multipliers. The only effect of this additional term is to restrict the original Lagrange multipliers to, $0 \le \alpha_i \le C$, instead of being simply non-negative. Lower values for C correspond to smaller penalties for outliers and a softer margin. The threshold b can still be found from Eq. (6), provided the corresponding Lagrange multiplier is not at the upper bound C.

Figures 3 and 4 illustrate the non-separable case, where the data points are the same as in Fig. 1, but with an outlier added. Figure 3 shows the results of using the standard form for computing the separating hyperplane ($C=\infty$) with a linear kernel. Notice the effect that the outlier had on the computed boundary and that nearly all the data points are support vectors. If we allow a softer margin and choose $C=10$, then Fig. 4 results, whose separating hyperplane is closer to that of Fig. 1. In this

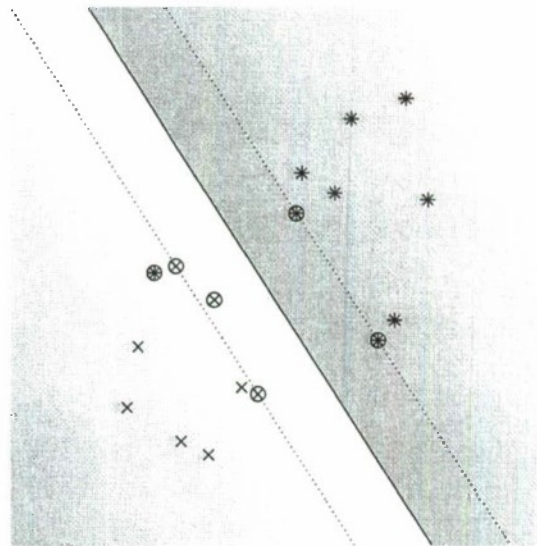Fig. 3: Linearly non-separable data with $C=\infty$.

Fig. 4: Linearly non-separable data with $C=10$.

case, the effect of the outlier has been minimized and there are fewer support vectors required to define the boundary. The kernel and the parameter C are the only items that the user needs to choose to run the SVM algorithm.

Fast algorithm

A problem with the baseline SVM algorithm can be seen by examining the objective function L_D in Eq. (8). In order to solve the quadratic programming (QP) problem entailed in finding the Lagrange multipliers that maximize L_D, one needs to compute the NxN Hessian matrix $(H_{ij}=y_iy_jK(x_i,x_j))$, where N is the number of training vectors. This matrix can require a significant amount of memory when N gets large (some problems have $N{\sim}10^5$) and performing operations on such a large matrix can be problematic.

Chunking[10] was one of the first methods used to increase the size of the problem that can be solved reasonably with the SVM algorithm. One separates the support vectors from the non-support vectors through an iterative filtering process. One begins by optimizing a small subset of the input vectors. The remaining vectors are placed in a testing subset and compared against the optimal hyperplane. A portion of the test vectors that are misclassified are added to the training subset and those vectors in the original training subset that were not support vectors are removed and added to the test subset. The final step is an optimization using all the support vectors, where none of the test vectors are misclassified. However, when the number of support vectors also becomes large, then the chunking method will no longer be adequate.

The next method discovered was to solve the QP problem in pieces, optimizing only a subset of the Lagrange multipliers in Eq. (8) at a time, keeping the others fixed.[11] At each step, a set of multipliers are determined such that, if the optimal solution is computed, the resulting L_D in Eq. (8) increases. One iterates until no feasible set exists, the objective function is maximum and the optimal solution is found. This method does not require the computation of the full Hessian matrix at any one time.

When the subset size of the latter algorithm is chosen to be two, it was found[5,12] that the optimization step could be performed analytically. This choice then removes the necessity of using numerical QP algorithms, which can be somewhat difficult to implement and use, and leads to a fast algorithm. At each step of this algorithm, a pair of Lagrange multipliers is optimized

such that the objective function L_D strictly increases. When no such pairs can be found the algorithm stops, having found the optimal solution. An advantage of the baseline algorithm is that the computation is independent of feature vector length. The Hessian matrix is computed once and then the optimization proceeds. This is not true for the fast implementation due to the multiple iterations. We have not tested the length dependence for the fast SVM. We implemented this algorithm in MATLAB© and compared the results with the baseline SVM algorithm.

RESULTS

Algorithms

To benchmark the classification performance of the SVM algorithm, we compared the results with the output of Multi-Layer Perceptron (MLP) and Radial Basis Function (RBF) neural networks from the MATLAB© Neural Network Toolbox.[13]

The MLP neural network configuration that we chose was a two layer network with 100 neurons in the hidden layer and was trained by back-propagation. We used a hyperbolic-tangent sigmoid transfer function for both layers. The MLP neural network was trained using the RPROP algorithm,[13,14] which was found to yield orders of magnitude faster training than the standard gradient descent method with momentum and adaptive learning rate. We used the default values of 1.2 and 0.5 for the RPROP increment and decrement parameters, respectively. The maximum weight change and initial value were set at 0.005 and the training error was chosen to be 1%, based on experiment to improve the classification. While there are a number of parameters to adjust in the RPROP-MLP neural network, the results are generally not very sensitive to the precise value of the parameters. We did notice an improvement of a few percent in the classification results when the maximum weight change was lowered from its default value of 50.

The RBF neural network, as implemented in MATLAB©, adds one neuron at a time to the network until the error goal is reached. The centers for the radial basis functions are selected from the training vectors and the width is chosen manually. During each iteration, the algorithm chooses the weights of the new neuron to be the feature vector that will result in the greatest decrease in total error. This leads to the smallest number of neurons, but takes a long time to train. We chose the width of the radial basis functions to be 25 and the training error to be 1%, both based on experiment to optimize the results.
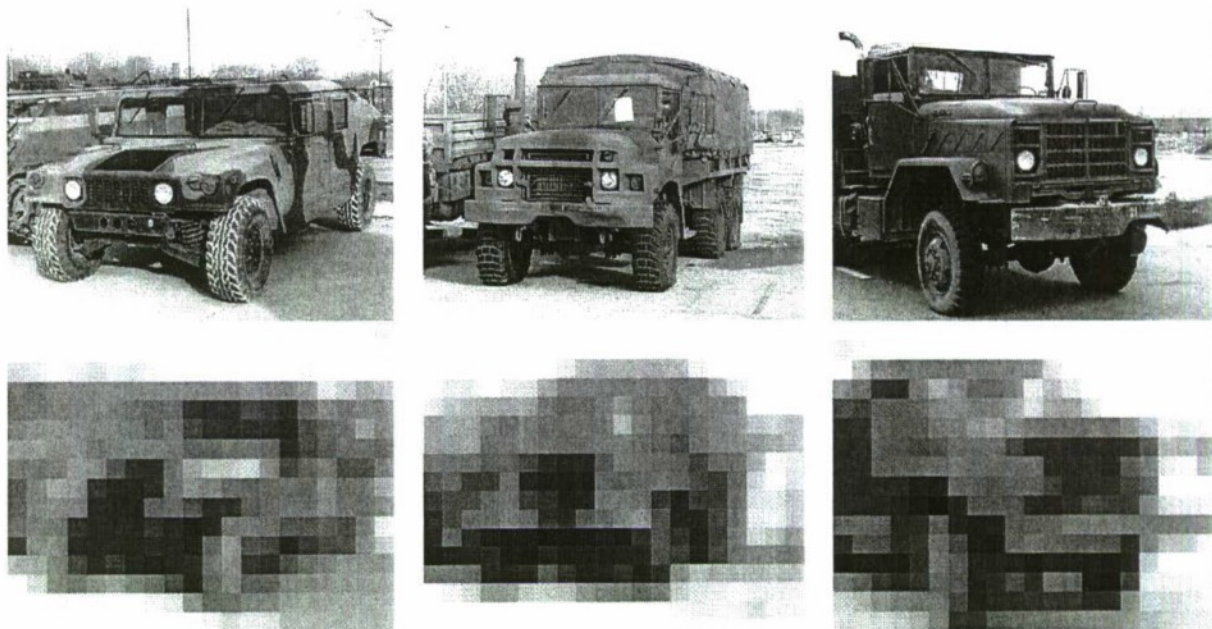
Fig. 5: Samples of vehicle imagery, original and reduced resolution.

The baseline and fast SVM algorithms were also implemented in MATLAB©. The former was developed by modifying and correcting an existing SVM toolbox[15] and the latter was coded directly from Refs. [5,12]. For both algorithms we chose to use a non-homogenous quadratic polynomial kernel and we set the non-separable parameter C to infinity for the baseline SVM and to 10 for the fast algorithm. Since the fast algorithm iterates over the Lagrange multipliers, it becomes quite slow when large values of C are used with non-separable data. The choice of C in this case was not critical due to the sparsity of the data. Since SVM is a binary classifier, it is necessary to train separate classifiers for each of the 5 vehicles. A test vector is then analyzed by each of the classifiers and the class with the largest output is chosen. The multi-class implementation of the SVM algorithm is a topic of active research.[5]

Military vehicle test

The first data set that we used to test the SVM algorithm consisted of images of military vehicles. We created 512x640 RGB digital images of five different vehicles (1 HMMWV and 4 trucks), taken from a variety of positions about the front of the vehicles with a digital camera. Since we wanted the algorithms to classify the vehicles based on their rough shapes, we converted the images to grayscale and reduced the resolution by a factor of 32 to 16x20. Examples of the imagery, before and after the resolution reduction, are shown in Fig. 5 and have been resized for display purposes. The imagery was

then converted column-wise to a vector. We normalized the resulting feature vectors to zero mean and unit standard deviation in order to remove, as much as possible, any extraneous features of the image classes. Since we had a limited number of images (50) for each vehicle, we resorted to a jackknife approach to training. We randomly chose N out of 50 images to train with and then tested with the remainder (except for the case when $N=49$, where we tested on each of the 50 images). We averaged over 50 iterations to obtain our results with each classifier being presented the same training and testing data.
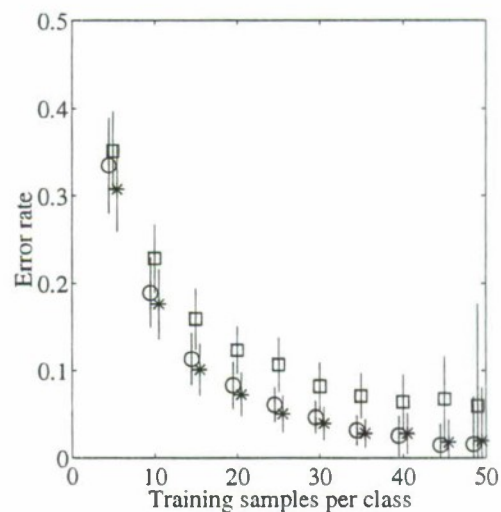

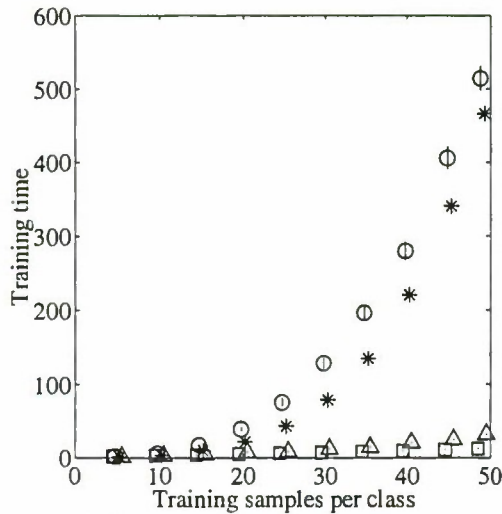Fig. 6: Error rate vs. number of training samples (□=MLP, O=RBF, *=SVM).

291

Fig. 7: Training time vs. number of training samples (□=MLP, O=RBF, *= baseline SVM, Δ=fast SVM).

The classification error rates for the three algorithms are displayed in Fig. 6 as a function of N, the number of training samples in each class. The plots are offset slightly for display clarity. The SVM error rates are slightly better (~1%) than the RBF results for the sparser training sets ($N<40$). The MLP error rates were significantly higher than the other two classifiers. We do not separate the baseline and fast SVM algorithm results since they find identical optimal solutions.

Vehicle	1	2	3	4	5
1	0.96	0.00	0.02	0.00	0.00
2	0.00	1.00	0.00	0.00	0.00
3	0.04	0.00	0.96	0.00	0.00
4	0.00	0.00	0.00	1.00	0.02
5	0.00	0.00	0.02	0.00	0.98

Table I: SVM classification results ($N=49$).

Figure 7 shows the training time as a function of N. The RPROP-MLP network trains substantially faster than the baseline SVM algorithm, which in turn is faster than the RBF network. However, the fast SVM algorithm is only slightly slower than the MLP network, while providing significantly better classification results. Small differences in training times should not be considered too seriously since there are many factors that affect computation time and there are many different approaches to MLP and RBF training. We have chosen one particular implementation and one particular set of parameters to test our algorithms, although we have sought for the best

parameters for this particular data set for each of the algorithms. Individual classification rates for the SVM algorithm is shown in Tables I for the case where 49 out of 50 feature vectors was used for training. The columns label the correct vehicle and the rows label the predicted vehicle. The HMMWV is vehicle 1 and the trucks are vehicles 2-5.

This turned out to be a relatively easy problem with fewer training examples than feature vector elements. The classification rates and training times for the two neural network configurations was somewhat of a surprise, since we did not expect them to perform well, due to the large feature vector size. A more complex and robust target data set is needed to fully differentiate between the three learning algorithms. An important addition to our future data collection will be different vehicles of the same class in order to measure inter-class variability.

Character recognition test

Due to the limited size of the previous database, we wanted to test our implementation of the algorithms against a larger set of data. Our second data set was taken from a MNIST database[6] of handwritten digits, 0-9, which consists of a 60,000 image training set and a 10,000 image testing set. The database mixes two separate databases from NIST, one using Census Bureau employees and the other using high school students. There are approximately 250 different writers in the training set, who are different from the writers in the testing set. Samples of the images are shown in Fig. 8. Other classification tests on this data set have
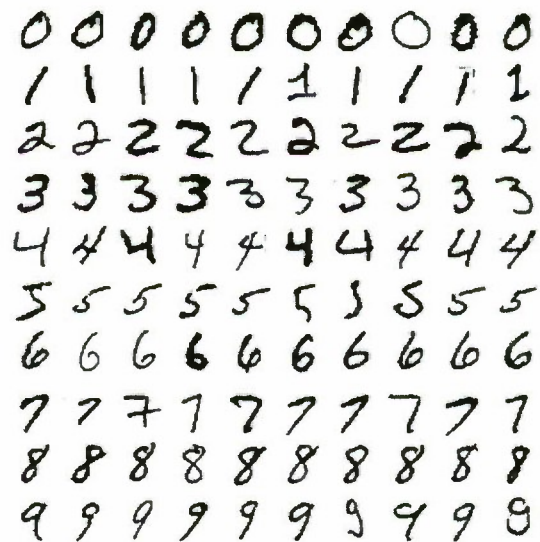

Fig. 8: Samples of MNIST digits.

292

Digit	0	1	2	3	4	5	6	7	8	9
0	97.9	0.0	0.5	0.1	0.2	0.8	0.6	0.0	0.3	0.5
1	0.0	98.9	0.3	0.0	0.1	0.0	0.2	1.6	0.2	0.5
2	0.2	0.3	95.3	0.9	0.0	0.1	0.1	1.2	0.1	0.2
3	0.1	0.1	0.9	95.5	0.0	1.6	0.0	0.2	1.3	0.7
4	0.0	0.0	0.4	0.0	96.2	0.2	0.9	0.4	0.4	1.3
5	0.3	0.1	0.0	0.9	0.0	95.1	0.8	0.1	0.1	0.2
6	0.6	0.3	0.5	0.0	0.8	0.9	96.3	0.0	0.6	0.3
7	0.1	0.1	1.0	0.8	0.1	0.1	0.1	94.7	0.3	1.2
8	0.6	0.3	0.9	1.3	0.2	0.7	0.8	0.2	95.3	0.4
9	0.2	0.0	0.3	0.5	2.3	0.6	0.0	1.7	0.5	94.7

Table 11: SVM classification table for the MNIST data.

demonstrated the superiority of the SVM algorithm over the MLP and RBF networks.[6]

The relatively slow training speed of the RBF and baseline SVM algorithms for large training sets led us to apply only the MLP and fast SVM algorithms to this data. For these cases, we simply input the entire (28x28) image into the algorithms as a feature vector. We trained the MLP network in two different ways, one being a straight 10-class implementation and the other being 10 binary classifiers, similar to the SVM setup. In the latter case, we included more samples of the target class in each training set than would normally occur. We found that this increases the classification performance.

The results of using 3,000 training vectors is shown in Table 11 for the fast SVM algorithm, which resulted in a total classification rate of 96.0% and a training time of 19.3 hours. Again, the rows label the predicted character and the columns label the correct character. The results for the two MLP networks were 92.0% classification rate in 0.6 hours for the 10-class implementation and 92.0% classification in 7.8 hours for the binary classifier. Based on previous tests, we expected the MLP binary classifier to outperform the 10-class algorithm. The fact that they achieved identical classification rates is somewhat of a surprise. As expected, the SVM classifier performed better than the MLP, although taking longer to train. We expect that SVM would still give better classification rates even for equal training times. Future test will examine this issue, especially in light of the training times and performance of the 10-class MLP network. The main point of this test was to show that, given sufficient training time, very large classification problems can be handled by the fast SVM algorithm, while retaining excellent classification results.

CONCLUSION

In summary, we have applied a relatively new learning algorithm, the Support Vector Machine (SVM), to a military vehicle classification problem and to a larger character recognition problem. Results were compared to standard neural network approaches. We demonstrated good classification results on the data set of military vehicle imagery for all the classifiers. The SVM algorithm performed slightly better than the Radial Basis Function (RBF) neural network, primarily on the sparser training sets. Both achieved significantly lower classification error rates than the Multi-Layer Perceptron (MLP) neural network. The fast implementation of SVM was much faster than the RBF network and only slightly slower than the MLP network, while still retaining the same classification performance of the baseline algorithm. If one were to compare classification rate as a function of training time, it would be seen that the fast SVM algorithm still performs significantly better than the MLP network for a given amount of training time.

The tests using the larger character recognition data set showed that the fast SVM algorithm could achieve better classification results with more than twice the training time as the MLP network. We expect that if data is taken for classification rate as a function of training time, the SVM algorithm will continue to outperform the MLP. The RBF network was not tested due to the extremely long training times that would be required. It has also been demonstrated that the fast SVM algorithm is capable of working with large training sets.

The advantages of the fast SVM implementation are that the solutions are unique, there are few parameters to adjust, it is outlier resistant, and it is a wide margin classifier, which gives good results even for sparse data sets. The baseline SVM algorithm has the further advantage that the training speed is independent of the feature vector length. We have not investigated the length dependence of the fast algorithm, which may in fact be only slightly higher than in the baseline algorithm. One disadvantage of SVM is that the testing phase can be slow if there are a large number of support vectors. The discovery of the reduced set method,[5,16] where a large number of support vectors are replaced by an effective vector, has already provided a substantial increase in the speed of the SVM test phase.

Future work consists of implementing improvements to the training and testing speeds, investigating methods to find the optimum kernel, and determining the usefulness of preprocessing the data. In particular, we are implementing a variant of the fast

algorithm that optimizes three multipliers simultaneously. We are also compiling larger and more complex data sets with which to test future implementations.

REFERENCES

1) R.E. Karlsen et. al, "Comparative study of wavelet methods for ground vehicle signature analysis," *Wavelet Applications III*, ed. H.H. Szu, SPIE Proc. **2762**, 314-24 (1996).

2) D.J. Gorsich, R.E. Karlsen, G.R. Gerhart and M.G. Genton, "Target versus background characterization: Second generation wavelets and support vector machines," *Targets and Backgrounds: Characterization and Representation V*, eds. W.R. Watkins, D. Clement and W.R. Reynolds, SPIE Proc. **3699**, 185-96 (1999).

3) V. Vapnik, *Statistical Learning Theory*, John Wiley and Sons, New York (1998).

4) C.J.C. Burges, "A tutorial on support vector machines for pattern recognition," *Data Mining and Knowledge Discovery* **2**, 1-47 (1998).

5) B. Schölkopf, C.J.C. Burges, and A.J. Smola, eds., *Advances in Kernel Methods, Support Vector Learning*, MIT Press, Cambridge MA (1999).

6) Y. Lecun et al., "Comparison of learning algorithms for handwritten digit recognition," *Proc. Int. Conf. Artificial Neural Networks* **II**, eds. F. Fogelman-Soulié and P. Gallinari, 53-60 (1995).

7) D.G. Luenberger, *Linear and Nonlinear Programming*, Addison-Wesley, Reading MA (1984).

8) R. Fletcher, *Practical Methods of Optimization*, John Wiley and Sons, Chichester (1987).

9) D.P. Bertsekas, *Nonlinear Programming*, Athena Scientific, Belmont MA (1995).

10) V. Vapnik, *Estimation of Dependences Based on Empirical Data*, Springer-Verlarg (1982).

11) E. Osuna, R. Freund and F. Girosi, "An improved training algorithm for support vector machines," *Proc. IEEE Neural Networks for Signal Processing* **VII**, eds. J. Principe et. al, 276-85 (1997).

12) J.C. Platt, "Sequential Minimal Optimization: A fast algorithm for training support vector machines," Tech Report MSR-TR-98-14, Microsoft Research (1998).

13) H. Demuth and M. Beale, *Neural Network Toolbox User's Guide*, The Mathworks Inc., Natick MA (1998).

14) M. Riedmiller and H. Braun, "A direct adaptive method for faster backpropagation learning: The RPROP algorithm," *Proc. IEEE Int. Conf. Neural Networks*, San Francisco (1993).

15) S. Gunn, "Support vector machines for classification and regression," Technical Report, Image Speech and Intelligent Systems Group, Univ. of Southampton (1998).

16) C.J.C. Burges, "Simplified support vector decision rules," *Proc. 13th Intl. Conf. on Machine Learning*, ed. L. Saitta, 71-77 (1996).

DISTRIBUTION LIST

FOR

TENTH ANNUAL
GROUND TARGET MODELING
& VALIDATION CONFERENCE

Dist. - A

ANDERSON, ALAN A
TELEDYNE BROWN ENGINEERING
1150 ACADEMY PARK LOOP SUITE 101
COLORADO SPRINGS, CO 80910-3715

AVERY, TERRANCE
US ARMY TACOM
AMSTA-TR-R, TVP TEAM
WARREN, MI 48397-5000

BENNETT, JOHN G
US ARMY TACOM
ATTN.: AMSTA-TR-S MS 263
WARREN, MI 48397-5000

BLASBAND, CHRIS
SURFACE OPTICS CORPORATION
11555 RANCHO BERNARDO ROAD
SAN DIEGO, CA 92127

BROWN, SCOTT D
ROCHESTER INST OF TECHNOLOGY
54 LOMB MEMORIAL DRIVE
ROCHESTER, NY 14623-5604

BRYK, DARRYL C
US ARMY TARDEC
ATTN.: AMSTA-TR-R MS 263
WARREN, MI 48397-5000

CAITO, STEVEN
US ARMY TACOM
ATTN: AMSTA-TR-S
WARREN, MI 48397-5000

CANNON, MELISA B
US ARMY AMCOM
AMSAM-RD-MG-IR
REDSTONE ARSENAL, AL 35898

CLINE, G. FRANCIS
SCIENTIFIC RESEARCH
CORPORATION
6767 OLD MADISON PIKE
HUNSTVILLE, AL 35806

CURRAN, ALLEN
THERMOANALYTICS, INC.
PO BOX 66
CALUMET, MI 49913

DENNEN, KEVIN
ERC INC.
BUILDING 4500
REDSTONE ARSENAL, AL 35898

DERENIAK, EUSTACE
UNIVERSITY OF ARIZONA
OPTICAL SCIENCES CTR
TUCSON, AZ 85721

ELLIS, DARIN
WAYNE STATE UNIVERSITY
87 EAST FERRY
DETROIT, MI 48202

EVANS, ROGER
US ARMY TACOM
AMSTA-TR-R MS 263
WARREN, MI 48397-5000

FREDRICK, DAVID
US ARMY TACOM
AMSTA-TR-R MS 263
WARREN, MI 48397-5000

FREESE, DOUG
US ARMY TACOM
Attn: AMSTA-JCH, Hit Avoidance BR
WARREN, MI 48397-5000

GAYNOR, WINN
DCS CORP.
1330 BRADDOCK DR.
ALEXANDRIA, VA 22314

GEOHAGAN, CLIFFORD
SAIC
6725 ODYSSEY DRIVE
HUNTSVILLE, AL 35806-3301

GERHART, GRANT
US ARMY TARDEC
ATTN: AMSTA-TR-R MS 263
WARREN, MI 48397-5000

GOETZ, RICHARD
US ARMY TACOM / TARDEC
ATTN: AMSTA-TR-R MS 263
WARREN, MI 48397-5000

GONDA, TERESA
US ARMY TACOM
ATTN: AMSTA-TR-R MS 263
WARREN, MI 48397-5000

HALSTEAD, DANIEL
ANGLE INC.
7406 ALBAN STATION COURT #A112
SPRINGFIELD, VA 22150

HUTCHINSON, HAROLD R
US ARMY TACOM
Attn: AMSTA-CM-XSF
WARREN, MI 48397-5000

JACKSON, WILLIAM
US ARMY TARDEC
AMSTA-TR-R SURVIVABILITY CENTER
MS 263
WARREN, MI 48397-5000

JACOBS, STEPHEN
AMHERST SYSTEMS INC
30 WILSON ROAD
BUFFALO, NY 14221

JACOBS, PAMELA
NIGHT VISION & ELEC SENSOR DIR
10221 BURBECK RD, STE 430 AMSEL-
RD-NV-CD-NVESD
FT. BELVOIR, VA 22060-5677

JAFOLLA, JAMES
SURFACE OPTICS CORPORATION
11555 RANCHO BERNARDO ROAD
SAN DIEGO, CA 92127-1441

JANICKI, PHILLIP
SIGNATURE RESEARCH INC.
PO BOX 344
OXFORD, MI 48371

JOHNSON, KEITH
THERMOANALYTICS, INC.
PO BOX 66
CALUMET, MI 49913

JOYNER, TOM
PRB ASSOCIATES INC
48150 SHAW ROAD BLDG 2109 STE
W233 UNIT 5
PATUXENT RIVER, MD 20670-1907

KARLSEN, ROBERT
US ARMY TACOM
ATTN: AMSTA-TR-R MS 263
WARREN, MI 48397-5000

LAUGHERY, SEAN
US ARMY TARDEC
AMSTA-TR-R MS 263
WARREN, MI 48310

LEONARD, DON
SVERDRUP TECHNOLOGY
4200 COLONEL GLENN HWY.
BEAVERCREEK, OH 45431

LESS, DAVID
THERMOANALYTICS INC.
PO BOX 66
CALUMET, MI 49913

LINDELL, MARTIN
GENERAL DYNAMICS LAND SYSTEMS
38500 MOUND ROAD
WARREN, MI 48310

LIVINGSTON, RAYMOND
TELEDYNE BROWN ENGINEERING
300 SPARKMAN DRIVE NW
HUNTSVILLE, AL 35807

MAKAR, ROBERT J.
AMHERST SYSTEMS INC
30 WILSON ROAD
BUFFALO, NY 14221

MCKEE, DOUGLAS
AMHERST SYSTEMS INC
30 WILSON ROAD
BUFFALO, NY 14221

MCKHEEN, TONY
US ARMY TARDEC
AMSTA-TR-R MS 263
WARREN, MI 48397

MCMANAMEY, JAMES R
NIGHT VISION & ELEC SENSOR DIR
10221 BURBECK RD, STE 430 AMSEL-
RD-NV-ST-VMS
FT. BELVOIR, VA 22060-5806

MEHTA, SUNIL
INGALLS SHIPBUILDING
PO BOX 149 MS 1090-11
PASCAGOULA, MS 39568

MILLER, DOUGLAS
US ARMY TACOM
AMSTA-TR-N 267
WARREN, MI 48397-5000

MOORE, RICHARD
ANGLE INC.
'406 ALBAN STATION COURT #A-112
SPRINGFIELD, VA 22150

MUENCH, PAUL
US ARMY TACOM
ATTN.: AMSTA-TR-R
WARREN, MI 48397

NEWHOUSE, JIM
LOCKHEED MARTIN SKUNK WORKS
1011 LOCKHEED WAY DEPT. 27-52
B/611 P/10
PALMDALE, CA 93099

PARKS, JACK
US ARMY TACOM
ATTN: AMSTA-TR-S
WARREN, MI 48397-5000

PARNELL, GREGORY
BOEING INFORMATION AND DEFENSE
SYS
PO BOX 3707, MC 4X-55
ELECROMAGNETICS DEPT

PASIK, MARK
GENERAL DYNAMICS LAND SYSTEMS
38500 MOUND ROAD MZ 436-20-26
STERLING HEIGHTS, MI 48310

PASSMORE, RONALD
US ARMY AMCOM
ATTN: AMSAM-RD-MG-IR
REDSTONE ARSENAL, AL 35898

PATTY, STEPHANIE R
NICHOLS RESEARCH CORPORATION
4040 S MEMORIAL PKWY PO BOX
400002
HUNTSVILLE, AL 35815-1502

PILOTTE, STEPHANIE
US ARMY AMCOM
BUILDING 6260, SED AMSAM-RD-BA-AT
REDSTONE ARSENAL, AL 35898

REED, JACK
US ARMY TACOM, TARDEC
AMSTA-TR-R SOM TEAM MS 263
WARREN, MI 48397-5000

REES, DONALD
GENERAL DYNAMICS LAND SYSTEMS
38500 MOUND RD MZ 436-20-27
STERLING HEIGHTS, MI 48310

RENIUS, OTTO
RENTECH INC
1091 CHRISTIAN HILLS
ROCHESTER HILLS, MI 48309

RICHARDS, MICHAEL
NICHOLS RESEARCH CORPORATION
1130 EGLIN PKWY, SUITE A SHALIMAR
CENTER
SHALIMAR, FL 32579

RINALD, DAN
US ARMY NGIC
220 SEVENTH ST NE
CHARLOTTESVILLE, VA 22902-5396

ROGERS, CHRIS
DEPT. OF ARMY, TSMO,PM, ITTS
DIRECTOR, THREAT SIMULATOR
MGMT OFFICE ATTN: AMSTI-ITTS-ST
REDSTONE ARSENAL, AL 35898-7461

ROUSE, WILLIAM
US ARMY EDGEWOOD CHEM AND BIO
CTR
ATTN.: AMSSB-RRT-PA
ABERDEEN PROV. GRD, MD 21010

SANDERS, JEFFREY S
SIMULATION TECHNOLOGIES INC
3307 BOB WALLACE AVENUE STE 3
HUNTSVILLE, AL 35805

SAUTER, DAVID P
ARL, BATTLEFIELD ENVIRON
DIRCTORATE
ATTN: AMSRL-BE-W
WHITE SANDS MISSILE, NM 88002-

SCHAFER, STEVE
SVERDRUP TECHNOLOGY
4200 COLONEL GLENN HWY.
BEAVERCREEK, OH 45431

SLEDGE, JOHN
46 OG/OGML
104 CHEROKEE AVENUE
EGLIN AFB, FL 32542-5600

SNORRASON, MAGNUS
CHARLES RIVER ANALYTICS
725 CONCORD AVENUE
CAMBRIDGE, MA 02138

SOLA, MARCOS
ARMY RESEARCH LAB (ARL)
2800 POWDER MILL RD ATTN.: AMSRL-SE-SE
ADELPHI, MD 20783-1197

STRATTAN, BARRY
LOCKHEED MARTIN SKUNK WORKS
1011 LOCKHEED WAY DEPT. 27-52 B/611 P/10
PALMDALE, CA 93099

TAYLOR, STACIE B
NGIC
220 7TH STREET NE ATTN: IANG-TSC
CHARLOTTESVILLE, VA 22902

THOMAS, DAVID
US ARMY TACOM
ATTN: AMSTA-TR-R
WARREN, MI 48397-5000

THOMPSON, JAMES L.
US ARMY TACOM
ATTN: AMSTA-TR-S
WARREN, MI 48397-5000

WALKER, GRAYSON
GENERAL DYNAMICS AMPHIBIOUS SYSTEMS
991 ANNAPOLIS WAY
WOODBRIDGE, VA 22191

WALLACE, WILLIAM R
TRACOR AEROSPACE
6500 TRACOR LANE
AUSTIN, TX 78725

WEBER, BRUCE
US ARMY RESEARCH LAB
2800 POWDERMILL RD AMSRL-SE-SE
ADELPHI, MD 20783

WITUS, GARY
TURING ASSOCIATES INC
1392 HONEY RUN DRIVE
ANN ARBOR, MI 48103

WOLFE, GREGORY J.
US ARMY TACOM
ATTN: AMSTA-TR-S
WARREN, MI 48397-5000

YOO, SUNG
NORTHROP GRUMMAN
600 HICKS ROAD
ROLLING MEADOWS, IL 60068

ZEGEL, FERDINAND H
RADIAN INC
5845 RICHMOND HWY STE 725
ALEXANDRIA, VA 22303

CAOLA, MIKE
SOWERBY RSCH CENTRE, BRITISH AEROSPACE
FPC 267, PO BOX 5
FILTON BRISTOL BS34 7QW UK

GEATCHES, RACHEL
SOWERBY RSCH CENTRE, BRITISH AEROSPACE
OPTICS & LASERS FPC 267, PO BOX 5
FILTON BRISTOL BS34 7QW UK

GILMORE, MARILYN
DERA
ROOM G007, BUILDING A2
FARNBOROUGH HAMPSHIRE GU14 0LX UK

MALTZ, MASHA
BEN-GURION UNIVERSITY OF THE NEGEV
DEPT OF INDUSTRIAL ENG AND MGMT PO BOX 653

MITCHELL, ALISTAIR
DEFENCE RESEARCH AGENCY - CHERTSEY
RM 204 BLDG 114 CHOBHAM LANE
CHERTSEY SURREY KT16 0EE UK

RAIDT, URBAN
FORSCHUNGS INSTITUTE
FORSCHUNGS INSTITUTE SCHLOSS KRESSBACH
TUBINGEN D-72072 GERMANY

RAPANOTTI, JOHN
DREV, CANADA
2459 BOUL. PIE-XI NORTH
VAL-BELAIR QUEBEC G3J 1X5 CANADA

REPASI, ENDRE
FGAN-FOM RESEARCH INSTITUTE
EISENSTOCKSTRASSE 12 D 76275 ETTLINGEN
GERMANY